

UniABG: Unified Adversarial View Bridging and Graph Correspondence for Unsupervised Cross-View Geo-Localization

Cuiqun Chen¹, Qi Chen^{1, 3, 4}, Bin Yang^{2†}, Xingyi Zhang^{1†}

¹ School of Computer Science and Technology, Anhui University, China

² School of Computer Science, Wuhan University, Wuhan, China

³ Information Materials and Intelligent Sensing Laboratory of Anhui Province, Anhui University

⁴ the Key Laboratory of Intelligent Computing and Signal Processing of Ministry of Education, Anhui University

chencuiqun@ahu.edu.cn, cq@stu.ahu.edu.cn, yangbin_cv@whu.edu.cn, xyzhanghust@gmail.com

† Corresponding authors

Abstract

Cross-view geo-localization (CVGL) matches query images (*e.g.*, drone) to geographically corresponding opposite-view imagery (*e.g.*, satellite). While supervised methods achieve strong performance, their reliance on extensive pairwise annotations limits scalability. Unsupervised alternatives avoid annotation costs but suffer from noisy pseudo-labels due to intrinsic cross-view domain gaps. To address these limitations, we propose *UniABG*, a novel dual-stage unsupervised cross-view geo-localization framework integrating adversarial view bridging with graph-based correspondence calibration. Our approach first employs View-Aware Adversarial Bridging (VAAB) to model view-invariant features and enhance pseudo-label robustness. Subsequently, Heterogeneous Graph Filtering Calibration (HGFC) refines cross-view associations by constructing dual inter-view structure graphs, achieving reliable view correspondence. Extensive experiments demonstrate state-of-the-art unsupervised performance, showing that UniABG improves Satellite \rightarrow Drone AP by +10.63% on University-1652 and +16.73% on SUES-200, even surpassing supervised baselines.

Code — <https://github.com/chenqi142/UniABG>

Extended version — <https://arxiv.org/abs/2511.12054>

Introduction

Cross-view geo-localization (CVGL) addresses the critical task of determining geographic coordinates for ground-level or aerial query images by matching them against georeferenced satellite imagery. The proliferation of unmanned aerial vehicles (UAVs) has significantly expanded CVGL’s scope, enabling centimeter-level positioning essential for urban navigation, autonomous systems, and augmented reality. UAVs’ flexible low-altitude imaging capability renders them particularly advantageous for fine-grained localization in complex environments.

Nevertheless, CVGL faces inherent challenges stemming from extreme cross-modal viewpoint and appearance variations. Prior efforts mitigate these through geometry-aware architectures (Regmi and Shah 2019; Liu and Li 2019), semantic-guided embedding (Hu et al. 2018; Zhao et al.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

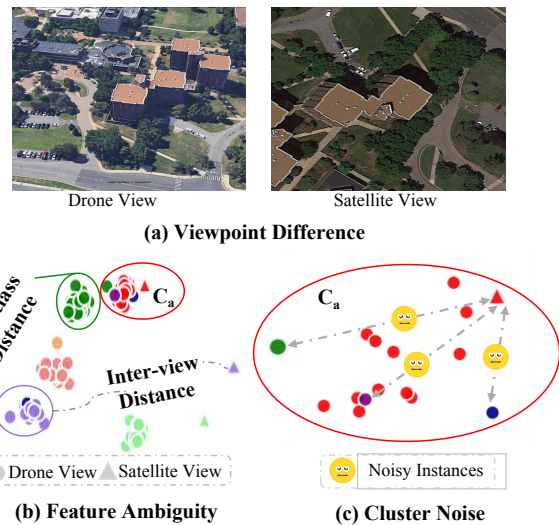


Figure 1: Key challenges in unsupervised CVGL. (a) Drastic appearance differences between drone and satellite views. (b) Feature space ambiguity, where the distance between different views of the same category may exceed the distance within different categories. (c) An enlarged view of the clustering space of category A (C_a). Noisy instances within clusters leading to incorrect pseudo-label association. **Different colors represent different categories.**

2024), and contrastive learning frameworks (Cai et al. 2019; Deuser, Habel, and Oswald 2023). While achieving notable performance, such supervised methods rely extensively on large-scale manually annotated cross-view image pairs, which incur prohibitive labelling costs.

To circumvent this limitation, recent research has shifted toward unsupervised cross-view geo-localization (UCVGL). Pioneering works include Li et al. (Li, Qian, and Xia 2024), who synthesize projected ground images aligned with satellite views to generate pseudo-pairs, and self-supervised approaches (Li et al. 2024; Wang et al. 2025a) that leverage foundation models with Expectation-Maximization (EM) mechanisms and consistency regularization for pseudo-label assignment. Wang et al. (Wang et al. 2025b; Yang

et al. 2022a, 2024) further adopt clustering-based contrastive learning to extract instance- and centroid-level features without supervision.

Despite these advances, the current UCVGL paradigm performs cross-view label association directly during training, overlooking two critical limitations. As shown in Fig. 1(a-b), divergent imaging altitudes and resolutions between drone/satellite modalities cause identical categories to exhibit larger feature-space distances across views than different categories within the same view. This view distribution misalignment propagates clustering errors through erroneous label associations. Furthermore, existing UCVGL methods employ simplistic label transfer mechanisms that fail under cluster impurity. Fig. 1(c) demonstrates how noisy instances trigger catastrophic association errors when directly mapped across views.

To address these challenges, we propose **UniABG**, a dual-stage framework for unsupervised cross-view geo-localization. Our approach synergizes adversarial view bridging with graph-based correspondence calibration, mitigating feature-level view gaps while establishing noise-robust cross-view correspondences. Through this dual-stage design, UniABG concurrently generates high-fidelity association pairs and learns view-consistent representations, achieving significant performance gains on standard unsupervised CVGL benchmarks.

In the first stage, to bridge the inherent geometric and spectral divergence between drone and satellite views, we establish a unified embedding space via View-Aware Adversarial Bridging (VAAB). This framework deploys feature extractors that compete adversarially against a domain classifier. By deliberately confusing the classifier across all three views – drone, satellite, and an Auxiliary Pseudo View (APV) – this optimization eliminates view-specific artifacts and forces discriminative, view-invariant feature learning. The APV is generated through cross-view style transfer between existing drone/satellite data. This synthetic view preserves structural semantics while simulating perspective transitions from low-altitude drone imagery to nadir-aligned satellite views, serving as a geometric intermediary that enables progressive viewpoint adaptation.

The second stage addresses error propagation caused by distorted neighborhood graph filtering structures during early clustering. Our heterogeneous graph filtering calibration (HGFC) module constructs dual graphs encoding inter-view structures to achieve robust correspondence matching. Specifically, for correspondence purification, we introduce mutual k-reciprocal neighbour filtering, which requires satellite candidates to neighbour both the drone image and its APV in the heterogeneous graphs. This symmetric constraint preserves geometrically consistent pairs while eliminating ambiguous matches from noisy clusters, thereby resolving the problematic noisy instances.

In summary, our principal contributions encompass:

(1) We propose the first dual-stage framework, *e.g.*, UniABG, for unsupervised cross-view geo-localization that integrates adversarial learning with graph correspondence filtering, jointly addressing the critical challenges of view discrepancy and association noise.

(2) We introduce a view-aware adversarial bridging strategy to mitigate cross-view discrepancies. This approach combines an auxiliary pseudo-view to model view-invariant embedding spaces while simultaneously enhancing cross-view feature discriminability.

(3) We develop a robust correspondence calibration mechanism through heterogeneous graph construction and mutual k-reciprocal neighbour filtering, significantly enhancing matching accuracy by eliminating ambiguous cross-view associations while preserving geometrically consistent pairs.

(4) We demonstrate state-of-the-art performance on University-1652 and SUES-200 benchmarks, where our approach outperforms all existing unsupervised methods and surpasses most supervised baselines, establishing new standards for unsupervised cross-view geo-localization.

Related Work

Supervised Cross-View Geo-localization

Cross-View Geo-Localization (CVGL) is critical for tasks such as autonomous navigation and augmented reality, where large viewpoint differences between ground and aerial imagery pose significant challenges. Early works focused on global feature learning, such as CVM-Net (Tian, Chen, and Shah 2017) and spatial-aware aggregation (Shi et al. 2019), while orientation priors (Liu and Li 2019) helped reduce spatial ambiguity. Subsequent studies addressed geometric misalignment, including Optimal Feature Transport (OFT) (Shi et al. 2020) and disentangled geometric learning (Zhang et al. 2023). Recent methods emphasize semantic reasoning and attention-based models to improve robustness. Geo-Net (Zhu et al. 2021) and SemGeo (Rodrigues and Tani 2023) leverage semantic cues, while Transformer-based approaches such as TransGeo (Zhu, Shah, and Chen 2022) and layer-wise correlation (Yang, Lu, and Zhu 2021) enhance long-range dependencies. Other strategies incorporate part-level correlation (Wang et al. 2021) and multi-candidate matching (Zhu, Yang, and Chen 2021) to improve performance.

Unsupervised Cross-View Geo-localization

The heavy reliance of supervised CVGL methods on labelled image pairs has motivated the development of unsupervised approaches. Li *et al.* (Li, Qian, and Xia 2024) proposed a framework that generates pseudo-labels through cross-view projection, while Li *et al.* (Li et al. 2024) introduced an EM-based self-supervised scheme but suffered from noisy label propagation. To alleviate such issues, Wang *et al.* (Wang et al. 2025b) employed clustering contrastive learning to learn cross-view representations without annotations, though it remained vulnerable to error accumulation. Building on these insights, we design a dual-stage framework that stabilizes clustering before performing cross-view association and integrates graph-based filtering to ensure contrastive supervision.

Method

This section presents UniABG, our proposed dual-stage framework for cross-view matching (Fig. 2). In the first

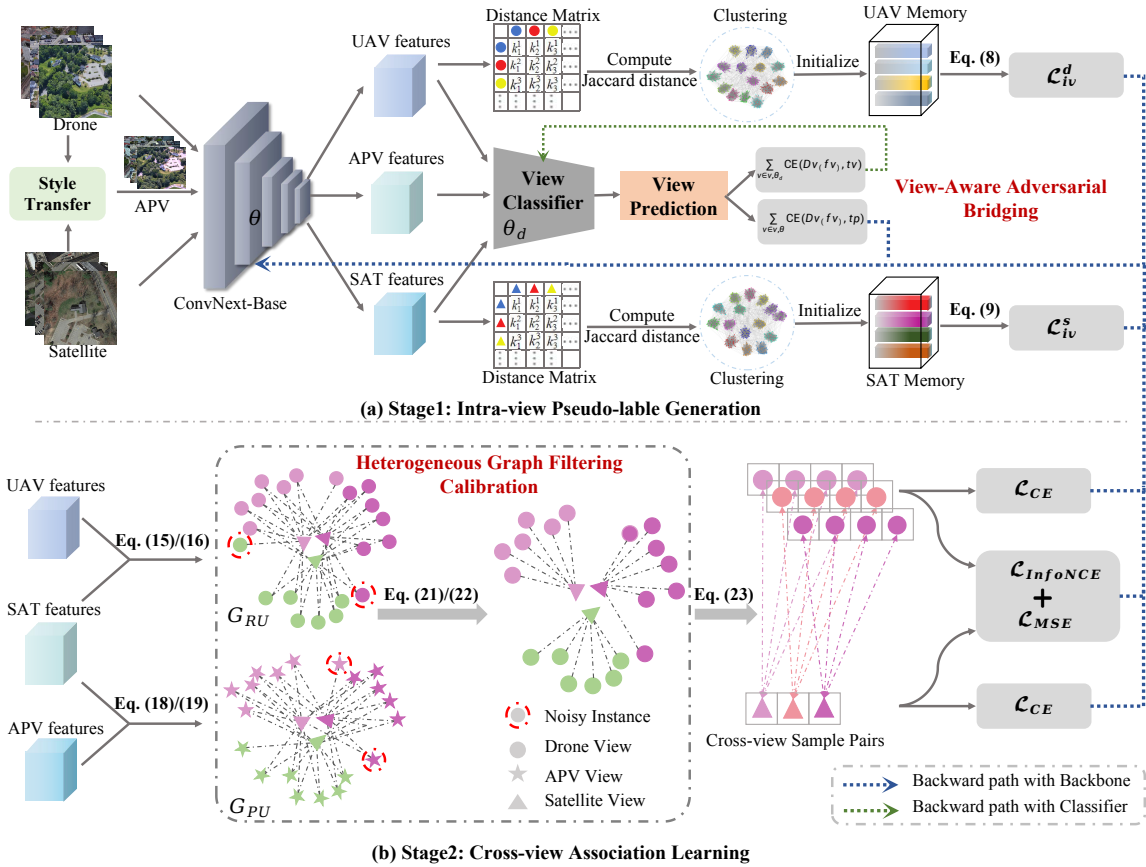


Figure 2: The overall architecture of our proposed UniABG is a dual-stage model. The first stage employs adversarial learning to reduce the differences between perspectives. The second stage constructs cross-view association data through heterogeneous graph filtering calibration for supervised learning (for more details, please refer to the text).

stage, pseudo-labels are generated via clustering and optimized through intra-view contrastive learning. These pseudo-labels are further refined using view-aware adversarial learning to bridge the cross-view discrepancy. Subsequently, heterogeneous graph filtering calibration refines cross-view correspondences, effectively enhancing matching reliability. UniABG achieves state-of-the-art performance on the University-1652 and SUES-200 datasets.

Problem Formulation. In the UCvGL, given an unlabeled drone-satellite image set, denoted as $\{X_d, X_s\}$, where $X^d = \{x_i^d | i = 1, 2, \dots, N^d\}$ represents the input drone image and $X^s = \{x_i^s | i = 1, 2, \dots, N^s\}$ represents the input satellite image. N^d and N^s represent the number of images in two views, respectively. The goal is to mine reliable positive and negative samples under this condition to complete the UCvGL task.

Dual-stage Unsupervised Cross-View Geo-Localization Baseline

Recent advances in cluster-based contrastive learning demonstrate significant potential for uncovering inherent data structures without labels. Building on this, we propose a dual-stage baseline for unsupervised cross-view geo-

localization. Our framework employs ConvNeXt-B (Xia et al. 2024) for feature extraction and operates as follows: (1) *Stage 1: Intra-view Pseudo-label Generation.* Clustering methods generate pseudo-labels within each view. To optimize feature discrimination, we apply intra-view contrastive learning constrained by a dual-contrastive strategy (Yang et al. 2022b). (2) *Stage 2: Cross-view Association Learning.* Cross-view data associations are constructed based on cosine similarity. These associations then drive cross-view contrastive learning to align view representations.

Intra-view Pseudo-label Generation. We first construct representative feature banks (memory dictionaries) by clustering unlabeled data within each view, serving as anchors for subsequent contrastive learning. Specifically, we extract features from drone (X^d) and satellite (X^s) images using a shared backbone network $\mathcal{F}_{\text{backbone}}$, yielding cross-view features

$$f_i^d = \mathcal{F}_{\text{backbone}}(x_i^d), i \in \{1, 2, \dots, N^d\}, \quad (1)$$

$$f_i^s = \mathcal{F}_{\text{backbone}}(x_i^s), i \in \{1, 2, \dots, N^s\}, \quad (2)$$

where $F_d = \{f_i^d\}_{i=1}^{N^d}$ and $F_s = \{f_i^s\}_{i=1}^{N^s}$ represent the feature sets of two views, respectively. To discover inherent data structures without labels, we apply DBSCAN (Es-

ter et al. 1996) separately on F_d and F_s to generate view-specific pseudo-labels

$$\hat{Y}^d = \text{DBSCAN}(F_d), \quad (3)$$

$$\hat{Y}^s = \text{DBSCAN}(F_s), \quad (4)$$

where $\hat{Y}^d = \{\hat{y}_1^d, \hat{y}_2^d, \dots, \hat{y}_{N^d}^d\}$ and $\hat{Y}^s = \{\hat{y}_1^s, \hat{y}_2^s, \dots, \hat{y}_{N^s}^s\}$ denote the pseudo labels of the input images X^d and X^s . Using these pseudo-labels, we compute the cluster centroids (prototypes) ϕ_k^d for drone view and ϕ_l^s for satellite view as the mean feature within each cluster

$$\phi_k^d = \frac{1}{|\mathcal{H}_k^d|} \sum_{\mathbf{f}_n^d \in \mathcal{H}_k^d} \mathbf{f}_n^d, \quad \phi_l^s = \frac{1}{|\mathcal{H}_l^s|} \sum_{\mathbf{f}_n^s \in \mathcal{H}_l^s} \mathbf{f}_n^s, \quad (5)$$

where $\mathcal{H}_k^d = \{\mathbf{f}_n^d \mid \hat{y}_n^d = k\}$ is the set of drone features belonging to cluster k , $\mathcal{H}_l^s = \{\mathbf{f}_m^s \mid \hat{y}_m^s = l\}$ is the set of satellite features belonging to cluster l , and $|\cdot|$ denotes set cardinality. Finally, we initialize the drone memory dictionary \mathcal{M}_d and satellite memory dictionary \mathcal{M}_s with their respective cluster prototypes

$$\mathcal{M}_d \leftarrow \{\phi_1^d, \phi_2^d, \dots, \phi_K^d\}, \quad \mathcal{M}_s \leftarrow \{\phi_1^s, \phi_2^s, \dots, \phi_L^s\}. \quad (6)$$

To further enhance the quality of intra-view pseudo-label generation, we employ an intra-view contrastive loss as a regularization mechanism. This approach optimizes the memory space distribution structure by reducing intra-memory feature distances while increasing inter-memory separation. Given drone and satellite query features q^d and q^s , we compute the contrastive loss for drone view and satellite view by the following equations

$$\mathcal{L}_{iv} = \mathcal{L}_{iv}^d + \mathcal{L}_{iv}^s, \quad (7)$$

$$\mathcal{L}_{iv}^d = -\log \frac{\exp(q^d \cdot \phi_+^d / \tau)}{\sum_{k=0}^K \exp(q^d \cdot \phi_k^d / \tau)}, \quad (8)$$

$$\mathcal{L}_{iv}^s = -\log \frac{\exp(q^s \cdot \phi_+^s / \tau)}{\sum_{l=0}^L \exp(q^s \cdot \phi_l^s / \tau)}, \quad (9)$$

where ϕ_+ is the positive memory corresponding to the pseudo label of q and the τ is a temperature. ϕ_k^d and ϕ_l^s represent the view cluster center features, respectively.

Cross-view Association Learning. Following the initial feature extraction, we perform the cross-view association learning for view matching. The objective here is to establish direct correspondences between individual drone and satellite images using generated pseudo-labels for supervised contrastive learning.

Given the feature sets for the drone view, $F_d = \{f_i^d\}_{i=1}^{N^d}$, and the satellite view, $F_s = \{f_j^s\}_{j=1}^{N^s}$, we construct cross-view associations using a greedy nearest-neighbour strategy. For each drone instance feature f_i^d , we compute its cosine similarity with every satellite instance feature in the set F_s . The satellite instance that yields the highest similarity is designated as the positive pair for f_i^d . This process identifies a

corresponding satellite instance index j^* for each drone instance i

$$j^* = \underset{j \in \{1, \dots, N^s\}}{\operatorname{argmax}} \left(\frac{(f_i^d)^\top f_j^s}{\|f_i^d\| \|f_j^s\|} \right). \quad (10)$$

Once the positive drone-satellite pair $(f_i^d, f_{j^*}^s)$ is identified via greedy matching (Eq. 10), all other satellite instances $\{f_j^s \mid j \neq j^*\}$ are treated as negatives for the drone query f_i^d . This allows us to define a cross-view contrastive learning objective: pull the positive pair closer in the embedding space while pushing negatives apart. We optimize the following supervised loss (Xia et al. 2024):

$$\mathcal{L}_{\text{sup}} = \mathcal{L}_{\text{InfoNCE}} + \mathcal{L}_{\text{MSE}} + \mathcal{L}_{\text{CE}}, \quad (11)$$

where $\mathcal{L}_{\text{InfoNCE}}$ learns scene-discriminative features, \mathcal{L}_{MSE} establishes cross-view correspondences, and \mathcal{L}_{CE} optimizes positive sample representations. However, this approach is critically vulnerable to noise in the initial association. Greedy matching is highly susceptible to visual ambiguities, generating erroneous positive pairs. Propagating these incorrect pseudo-labels into the supervised learning framework leads to significant error accumulation, thereby fundamentally limiting model performance. This underscores the necessity of our bridging and correspondence purification strategy in *UniABG*, which refines pseudo-labels and enhances cross-view robustness.

View-Aware Adversarial Bridging

A fundamental challenge in unsupervised cross-view geolocalization (UCVGL) stems from the significant domain gap induced by drastic perspective differences between drone and satellite imagery. This discrepancy, characterised by extreme scale variations, geometric distortions, and illumination disparities, results in substantial misalignment of feature distributions across views (Liu, Ye, and Du 2024). *This misalignment is particularly detrimental in the unsupervised setting.* However, naive cross-view clustering assumes cross-view feature compatibility but fails to bridge domain gaps, yielding performance-limiting noisy pseudo-labels.

To explicitly mitigate this cross-view distribution divergence for robust clustering, we propose a view-aware adversarial bridging (VAAB) module. Inspired by cross-modal shared representation learning (Lin et al. 2022; Yang, Hu, and Hu 2025; Yang, Chen, and Ye 2023; Yang et al. 2023; Yang, Chen, and Ye 2024; Yao et al. 2025), VAAB employs a synthetic pseudo-view as a transitional domain with adversarial constraints. Its core mechanism uses triplet-view adversarial training to reduce view-specific characteristics, learning discriminative cross-view representations.

Auxiliary Pseudo View. Referring to (Reinhard et al. 2002; Wang et al. 2025b), we obtain the auxiliary pseudo view (APV) by performing a global color transfer from the satellite domain to the drone domain. Both images are first converted into the Lab color space, and then channel-wise statistics (mean μ_c and standard deviation σ_c for $c \in \{L, a, b\}$) are computed. For the satellite domain, we compute the global mean and standard deviation (μ_c^s, σ_c^s) for

each channel $c \in \{L, a, b\}$ over the entire dataset. For each input drone image, we compute its individual statistics (μ_c^d , σ_c^d) and apply the following channel-wise transformation

$$l'_c = \frac{\sigma_c^s}{\sigma_c^d}(l_c - \mu_c^d) + \mu_c^s, \quad (12)$$

where l_c and l'_c denote the original and transformed pixel values, and the superscripts s and d indicate satellite and drone domains, respectively. This transformation standardizes drone image appearance while preserving its structural content, forming the auxiliary pseudo view.

Adversarial View Bridging. While style transfer provides pixel-level approximation, it fails to ensure semantic-level feature consistency across views. To achieve deeper domain invariance essential for cross-view association, we introduce an adversarial view bridging strategy. This approach transcends conventional feature alignment by explicitly harmonizing feature distributions through adversarial training. The core mechanism employs a view discriminator D_v that attempts to classify feature origins (drone x^d , satellite x^s , or APV x^p). By adversarially training the backbone to confound D_v , we force it to suppress view-specific artefacts and extract viewpoint-invariant representations.

Formally, features from all views are extracted via shared backbone F_B :

$$f_i^d = F_B(x_i^d), \quad f_i^s = F_B(x_i^s), \quad f_i^p = F_B(x_i^p). \quad (13)$$

A view discriminator D_v is trained to identify perspective origins, while F_B adversarially learns to induce geometric confusion – making features indistinguishable across views yet discriminative for location. This is achieved through a unified adversarial objective

$$\mathcal{L}_{VAAB} = \sum_{v \in \mathcal{V}, \theta_d} \text{CE}(D_v(f^v), t^v) + \sum_{v \in \mathcal{V}, \theta} \text{CE}(D_v(f^v), t^p), \quad (14)$$

where $\mathcal{V} = d, s, p$ denotes drone/satellite/APV views. Thus, t^d , t^s , and t^p are the corresponding view labels for three views. θ_d denotes the parameters of the view classifier and θ denotes the parameters of the backbone. This convergence emerges from complementary roles where the first item of \mathcal{L}_{VAAB} forces the discriminator to explore view-specific patterns, while the second drives the backbone to eliminate spectral and geometric biases. This establishes a foundation for reliable pseudo-label generation in subsequent clustering and correspondence optimization stages.

Heterogeneous Graph Filtering Calibration

The quality of cross-view association data is crucial to the effectiveness of subsequent supervised training. However, traditional methods mostly rely solely on the feature similarity within a single view to perform matching. This approach is highly sensitive to noise and is prone to generating a large number of ambiguous matches, making it difficult to ensure the quality of the association data (Xu et al. 2025).

To address this, we propose a Heterogeneous Graph Filtering Calibration (HGFC) module based on the auxiliary pseudo-view. HGFC aims to construct high-quality cross-view association data by exploiting the structural consensus across two complementary views: drone-to-satellite and

APV-to-satellite. Our key hypothesis is that true matching relationships should exhibit structural consistency across multiple feature manifolds, thereby effectively filtering out incorrect matches and improving the accuracy and reliability of the associations.

Heterogeneous Graph Construction. Given the extracted feature sets of drone and satellite images, denoted as $F_d = \{f_i^d\}$ and $F_s = \{f_j^s\}$ respectively, we compute the cosine similarity between each drone feature and all satellite features

$$\text{sim}(f_i^d, f_j^s) = \frac{(f_i^d)^\top f_j^s}{\|f_i^d\| \cdot \|f_j^s\|}. \quad (15)$$

Then, for each f_i^d , we construct its k -nearest satellite neighbors

$$N_k^{RU}(f_i^d) = \text{Top-}k \text{ sim}(f_i^d, f_j^s)_{f_j^s \in F_s}, \quad (16)$$

which defines the edge set \mathcal{E}_{RU} of the *Real-to-Real Graph* $G_{RU} = (F_d, F_s, \mathcal{E}_{RU})$. The edge set \mathcal{E} is defined as the connection relationships based on feature similarity

$$\mathcal{E}_{RU} = \{(f^d, f^s) \mid f^s \in N_k(f^d)\}. \quad (17)$$

Similarly, we use the auxiliary pseudo view feature set $F_p = \{f_i^p\}$ and compute its similarity with F_s

$$\text{sim}(f_i^p, f_j^s) = \frac{(f_i^p)^\top f_j^s}{\|f_i^p\| \cdot \|f_j^s\|}, \quad (18)$$

and define its neighbour set as

$$N_k^{PU}(f_i^p) = \text{Top-}k \text{ sim}(f_i^p, f_j^s)_{f_j^s \in F_s}, \quad (19)$$

which forms the *Pseudo-to-Real Graph* $G_{PU} = (F_p, F_s, \mathcal{E}_{PU})$. The edge set \mathcal{E} is defined as the connection relationships based on feature similarity

$$\mathcal{E}_{PU} = \{(f^p, f^s) \mid f^s \in N_k(f^p)\}. \quad (20)$$

Topological Consistency Alignment. To filter out noisy matches, we align the graph structures across the two views. Specifically, we apply mutual k -nearest neighbour (mutual-KNN) verification and evaluate the consistency of neighbours between the two graphs. For a satellite feature f_j^s , let $N_{RU}(f_j^s)$ and $N_{PU}(f_j^s)$ denote the sets of UAV and pseudo-view features that link to f_j^s in the respective graphs. The cross-graph consistency score is computed as

$$s_{ij}^{\text{cross}} = \frac{|N_k^{RU}(f_j^s) \cap N_k^{PU}(f_j^s)|}{k}, \quad (21)$$

and only satellite features with $s_{ij}^{\text{cross}} > \tau$ are retained for reliable association.

Semantics-Guided Intra-Cluster Weighted Voting. To suppress outliers and leverage the collective evidence within each cluster, we introduce a semantically guided weighted voting mechanism that uses semantic similarity and structural confidence to refine pseudo-labels. For a candidate pair (f_i^d, f_j^s) , the weighted confidence is defined as

$$\omega_{ij} = \text{sim}(f_i^d, f_j^s) \cdot s_{ij}^{\text{cross}}. \quad (22)$$

Within each UAV cluster \mathcal{C}_k , the final association is obtained via

$$\hat{y}_i = \arg \max_c \sum_{f_j^s \in \mathcal{C}_c} \omega_{ij}, \quad (23)$$

where \hat{y}_i is the refined pseudo-label, and \mathcal{C}_c denotes the c -th cluster of satellite features. This heterogeneous graph filtering calibration ensures that associations are not only based on appearance similarity but also structurally validated across heterogeneous views, thus significantly improving the precision of cross-view matching under unsupervised settings.

The Overall Loss Function

UniABG employs a dual-stage optimization strategy. The first stage leverages view-aware adversarial learning to reduce cross-view discrepancies and extract discriminative features, facilitating robust clustering. The stage 1 objective combines

$$\mathcal{L}_{\text{stage1}} = \mathcal{L}_{\text{iv}} + \lambda \cdot \mathcal{L}_{\text{VAAB}}, \quad (24)$$

where \mathcal{L}_{iv} enforces intra-view feature consistency, and $\mathcal{L}_{\text{VAAB}}$ ensures cross-view domain invariance. These losses jointly learn discriminative, view-invariant representations.

Using heterogeneous graph filtering, we generate high-confidence matching pairs from the clustered view images. These purified pairs serve as supervisory signals for the second stage, optimized via supervised contrastive loss \mathcal{L}_{sup} .

Experiments

Datasets and Experimental Settings

We evaluate our method on two public cross-view geolocation benchmarks: University-1652 (Zheng, Wei, and Yang 2020) and SUES-200 (Zhu et al. 2023) using Recall@K (R@K) and Average Precision (AP).

University-1652 (Zheng, Wei, and Yang 2020) contains 1,652 buildings from 72 universities, featuring satellite, ground, and synthetic drone views to support cross-view retrieval and navigation.

SUES-200 (Zhu et al. 2023) focuses on altitude variations, providing 24,120 drone images across four heights (150m–300m) and 200 scenes for drone-satellite matching.

Implementation Details. We implement our framework in PyTorch and conduct training on four NVIDIA RTX 4090 GPUs (24GB each). All input images are resized to 384×384 pixels. We adopt ConvNeXt-Base and a supervised objective as baseline (Xia et al. 2024). Both training stages use a batch size of 24 (12 drone + 12 satellite image pairs per GPU) for 5 epochs. Optimization employs AdamW with an initial learning rate of $1e^{-3}$, following a cosine decay schedule. The hyperparameter λ is fixed at 0.1.

Ablation Study

We ablate two core components: View-Aware Adversarial Bridging (VAAB) and Heterogeneous Graph Filtering Calibration (HGFC). These progressively mitigate cross-view variations while enhancing association quality. Table 1 quantifies contributions under identical settings: 1) *Effect*

Method	Drone \rightarrow Satellite		Satellite \rightarrow Drone	
	R@1	AP	R@1	AP
B	35.94	41.64	65.47	35.61
B+VAAB	60.36 $\uparrow 24.42$	65.03 $\uparrow 23.57$	80.74 $\uparrow 15.27$	58.77 $\uparrow 23.16$
B+HGFC	90.83 $\uparrow 54.89$	92.85 $\uparrow 51.21$	94.57 $\uparrow 29.10$	90.87 $\uparrow 55.26$
B+HGFC+VAAB	93.62 $\uparrow 2.79$	94.61 $\uparrow 1.76$	95.43 $\uparrow 0.86$	93.29 $\uparrow 2.42$

Table 1: The results of ablation studies of the proposed method on the University-1652. B represents the baseline.




Figure 3: The t-SNE visualization of the baseline and our proposed method. We randomly selected 50 location categories from the University-1652 dataset, with each color representing a location. “ \triangle ” represents satellite features, and “ \circ ” represents UAV features. Red circles indicate ambiguous matches.

of the HGFC: HGFC enforces structural consistency for robust associations. For Drone \rightarrow Satellite, it boosts R@1 by +54.89% and AP by +51.21%. For Satellite \rightarrow Drone, gains reach +29.1% R@1 and +55.26% AP. This verifies that structural filtering is essential to overcome ambiguous appearance-based matches. 2) *Effect of the VAAB:* VAAB reduces domain gaps through adversarial pseudo-view (APV) generation. It yields further improvements: +2.79% R@1 / +1.76% AP for Drone \rightarrow Satellite and +0.86% R@1 / +2.42% AP for Satellite \rightarrow Drone. This confirms that explicit view-invariant learning via APV is critical for feature robustness. 3) *Visual Analysis:* Fig. 3 compares t-SNE visualizations of baseline and our UniABG. Baseline features show poor clustering with incorrect matches. Our model tightly clusters cross-view features of the same class while separating different classes with distinct margins.

Comparison with the State-of-the-art Methods

We evaluate our method on University-1652 and SUES-200 benchmarks. On University-1652 (Table 2), our approach advances the state-of-the-art by significant margins: +7.67% Drone \rightarrow Satellite R@1 and +10.63% Satellite \rightarrow Drone AP.

SUES-200 results (Table 3) further demonstrate superiority, with particularly notable gains at the challenging 150m altitude. For Drone \rightarrow Satellite retrieval, we achieve +15.5% R@1 and +9.0% AP. In the Satellite \rightarrow Drone direction, improvements reach +11.25% R@1 and +16.73% AP.

	Method	Drone → Satellite		Satellite → Drone	
		R@1	AP	R@1	AP
S	Zhen et al. (Zheng, Wei, and Yang 2020)	59.69	64.8	73.18	59.4
	PCL (Tian et al. 2021)	79.47	83.63	87.69	78.51
	F3-net (Sun, Liu, and Yuan 2023)	78.64	81.60	-	-
	Sample4Geo (Deuser, Habel, and Oswald 2023)	92.65	93.81	95.14	91.39
	DAC (Xia et al. 2024)	94.67	95.50	96.43	93.79
	QDFL (Hu et al. 2025)	95.00	95.83	97.15	94.57
U	Li et al. (Li et al. 2024)	70.29	74.93	79.03	61.03
	Wang et al. (Wang et al. 2025b)	85.95	90.33	94.01	82.66
	UniABG	93.62	94.61	95.43	93.29

Table 2: Performance (R@1 % and AP %) comparison with SOTA methods on the University-1652 dataset. Supervised (S); Unsupervised (U). Best results are marked in bold.

Drone → Satellite									
	Method	150m		200m		250m		300m	
		R@1	AP	R@1	AP	R@1	AP	R@1	AP
S	SUES-200 (Zhu et al. 2023)	55.65	61.92	66.78	71.55	72.00	76.43	74.05	78.26
	FSRA (Dai et al. 2021)	68.25	73.45	83.00	85.99	90.68	92.27	91.95	93.46
	MCCG (Shen et al. 2023)	82.22	85.47	89.38	91.41	93.82	95.04	95.07	96.20
	DAC (Xia et al. 2024)	96.80	97.54	97.48	97.97	98.20	98.62	97.58	98.14
	QDFL (Hu et al. 2025)	93.97	95.42	98.25	98.67	99.30	99.48	99.31	99.48
U	Wang et al. (Wang et al. 2025b)	76.90	84.95	87.88	92.60	92.98	95.66	95.10	96.92
	UniABG	92.40	93.95	97.32	97.92	98.07	98.55	98.67	98.98
Satellite → Drone									
	Method	150m		200m		250m		300m	
		R@1	AP	R@1	AP	R@1	AP	R@1	AP
S	SUES-200 (Zhu et al. 2023)	75.00	55.46	85.00	66.05	86.25	69.94	88.75	74.46
	FSRA (Dai et al. 2021)	83.75	76.67	90.00	85.34	93.75	90.17	95.00	92.03
	MCCG (Shen et al. 2023)	97.50	93.63	98.75	96.70	98.75	98.28	98.75	98.05
	DAC (Xia et al. 2024)	97.50	94.06	98.75	96.66	98.75	98.09	98.75	97.87
	QDFL (Hu et al. 2025)	98.75	95.10	98.75	97.92	100.00	99.07	100.00	99.07
U	Wang et al. (Wang et al. 2025b)	87.50	74.81	92.50	87.15	96.25	91.20	98.75	94.52
	UniABG	98.75	91.54	98.75	97.06	100.00	98.32	98.75	97.58

Table 3: Performance (R@1 % and AP %) comparison with SOTA methods on the SUES-200 dataset. Supervised (S); Unsupervised (U). Best results are marked in bold.

Conclusion

We propose UniABG, a dual-stage unsupervised framework addressing viewpoint discrepancy and association noise in cross-view geo-localization. Specifically, our approach integrates adversarial learning to align drone-satellite feature distributions, while simultaneously employing heterogeneous graph filtering to resolve ambiguous matches and suppress error propagation during association construction. Experimental results demonstrate state-of-the-art unsupervised performance on University-1652 and SUES-200 benchmarks. Crucially, UniABG establishes that combin-

ing structural graph filtering with adversarial learning generates robust feature representations for label-free localization. Collectively, this work provides both an effective solution for cross-view geo-localization and a promising paradigm for future cross-modal matching research.

Acknowledgments

This work is partially supported by National Natural Science Foundation of China under Grants (62306215, 62501428), Postdoctoral Fellowship Program of China Postdoctoral Science Foundation (GZC20241268, 2024M762479), Hubei

Postdoctoral Talent Introduction Program (2024HBB-HJD070) and Hubei Provincial Natural Science Foundation of China (2025AFB219). The numerical calculations in this paper had been supported by the super-computing system in the Supercomputing Center of Wuhan University.

References

- Cai, S.; Guo, Y.; Khan, S.; Hu, J.; and Wen, G. 2019. Ground-to-aerial image geo-localization with a hard exemplar reweighting triplet loss. In *Proceedings of the IEEE/CVF international conference on computer vision*, 8391–8400.
- Dai, M.; Hu, J.; Zhuang, J.; and Zheng, E. 2021. A transformer-based feature segmentation and region alignment method for UAV-view geo-localization. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(7): 4376–4389.
- Deuser, F.; Habel, K.; and Oswald, N. 2023. Sample4geo: Hard negative sampling for cross-view geo-localisation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 16847–16856.
- Ester, M.; Kriegel, H.-P.; Sander, J.; Xu, X.; et al. 1996. A density-based algorithm for discovering clusters in large spatial databases with noise. In *kdd*, volume 96, 226–231.
- Hu, S.; Feng, M.; Nguyen, R. M.; and Lee, G. H. 2018. Cvm-net: Cross-view matching network for image-based ground-to-aerial geo-localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7258–7267.
- Hu, S.; Shi, Z.; Jin, T.; and Liu, Y. 2025. Query-Driven Feature Learning for Cross-View Geo-Localization. *IEEE Transactions on Geoscience and Remote Sensing*, 63: 1–15.
- Li, G.; Qian, M.; and Xia, G.-S. 2024. Unleashing unlabeled data: A paradigm for cross-view geo-localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16719–16729.
- Li, H.; Xu, C.; Yang, W.; Yu, H.; and Xia, G.-S. 2024. Learning Cross-View Visual Geo-Localization Without Ground Truth. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–17.
- Lin, X.; Li, J.; Ma, Z.; Li, H.; Li, S.; Xu, K.; Lu, G.; and Zhang, D. 2022. Learning Modal-Invariant and Temporal-Memory for Video-based Visible-Infrared Person Re-Identification. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 20941–20950.
- Liu, F.; Ye, M.; and Du, B. 2024. Learning a generalizable re-identification model from unlabelled data with domain-agnostic expert. *Visual Intelligence*, 2(1): 28.
- Liu, L.; and Li, H. 2019. Lending orientation to neural networks for cross-view geo-localization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5624–5633.
- Regmi, K.; and Shah, M. 2019. Bridging the domain gap for ground-to-aerial image matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 470–479.
- Reinhard, E.; Adhikhmin, M.; Gooch, B.; and Shirley, P. 2002. Color transfer between images. *IEEE Computer graphics and applications*, 21(5): 34–41.
- Rodrigues, R.; and Tani, M. 2023. Semgeo: Semantic keywords for cross-view image geo-localization. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Shen, T.; Wei, Y.; Kang, L.; Wan, S.; and Yang, Y.-H. 2023. MCCG: A ConvNeXt-based multiple-classifier method for cross-view geo-localization. *IEEE Transactions on Circuits and Systems for Video Technology*, 34(3): 1456–1468.
- Shi, Y.; Liu, L.; Yu, X.; and Li, H. 2019. Spatial-aware feature aggregation for cross-view image based geo-localization. In *Proceedings of the 33rd Annual Conference on Neural Information Processing Systems, NeurIPS 2019*, 10090–10100.
- Shi, Y.; Yu, X.; Liu, L.; Zhang, T.; and Li, H. 2020. Optimal feature transport for cross-view image geo-localization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 11990–11997.
- Sun, B.; Liu, G.; and Yuan, Y. 2023. F3-Net: Multiview scene matching for drone-based geo-localization. *IEEE Transactions on Geoscience and Remote Sensing*, 61: 1–11.
- Tian, X.; Shao, J.; Ouyang, D.; and Shen, H. T. 2021. UAV-satellite view synthesis for cross-view geo-localization. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(7): 4804–4815.
- Tian, Y.; Chen, C.; and Shah, M. 2017. Cross-view image matching for geo-localization in urban environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3608–3616.
- Wang, T.; Zheng, Z.; Yan, C.; Zhang, J.; Sun, Y.; Zheng, B.; and Yang, Y. 2021. Each part matters: Local patterns facilitate cross-view geo-localization. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(2): 867–879.
- Wang, X.; Liu, L.; Yang, B.; Ye, M.; Wang, Z.; and Xu, X. 2025a. TokenMatcher: Diverse Tokens Matching for Unsupervised Visible-Infrared Person Re-Identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 7934–7942.
- Wang, X.; Zhang, L.; Fan, Z.; Liu, Y.; Chen, C.; and Deng, F. 2025b. From Coarse to Fine: A Matching and Alignment Framework for Unsupervised Cross-View Geo-Localization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 8024–8032.
- Xia, P.; Wan, Y.; Zheng, Z.; Zhang, Y.; and Deng, J. 2024. Enhancing Cross-View Geo-Localization With Domain Alignment and Scene Consistency. *IEEE Transactions on Circuits and Systems for Video Technology*, 34(12): 13271–13281.
- Xu, Y.; Wu, M.; Guo, Z.; Cao, M.; Ye, M.; and Laaksonen, J. 2025. Efficient text-to-video retrieval via multi-modal multi-tagger derived pre-screening. *Visual Intelligence*, 3(1): 1–13.
- Yang, B.; Chen, J.; Chen, C.; and Ye, M. 2024. Dual Consistency-Constrained Learning for Unsupervised

- Visible-Infrared Person Re-Identification. *IEEE Transactions on Information Forensics and Security*, 19: 1767–1779.
- Yang, B.; Chen, J.; Ma, X.; and Ye, M. 2023. Translation, Association and Augmentation: Learning Cross-Modality Re-Identification From Single-Modality Annotation. *IEEE Transactions on Image Processing*, 32: 5099–5113.
- Yang, B.; Chen, J.; and Ye, M. 2023. Towards Grand Unified Representation Learning for Unsupervised Visible-Infrared Person Re-Identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 11069–11079.
- Yang, B.; Chen, J.; and Ye, M. 2024. Shallow-Deep Collaborative Learning for Unsupervised Visible-Infrared Person Re-Identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16870–16879.
- Yang, B.; Ye, M.; Chen, J.; and Wu, Z. 2022a. Augmented Dual-Contrastive Aggregation Learning for Unsupervised Visible-Infrared Person Re-Identification. In *ACM MM*, 2843–2851.
- Yang, B.; Ye, M.; Chen, J.; and Wu, Z. 2022b. Augmented dual-contrastive aggregation learning for unsupervised visible-infrared person re-identification. In *Proceedings of the 30th ACM International Conference on Multimedia*, 2843–2851.
- Yang, H.; Lu, X.; and Zhu, Y. 2021. Cross-view geo-localization with layer-to-layer transformer. *Advances in Neural Information Processing Systems*, 34: 29009–29020.
- Yang, Y.; Hu, W.; and Hu, H. 2025. Progressive Cross-Modal Association Learning for Unsupervised Visible-Infrared Person Re-Identification. *IEEE Transactions on Information Forensics and Security*, 20: 1290–1304.
- Yao, H.; Yang, B.; Huang, W.; Du, B.; and Ye, M. 2025. Unsupervised Visible-Infrared Person Re-identification under Unpaired Settings. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11916–11926.
- Zhang, X.; Li, X.; Sultani, W.; Zhou, Y.; and Wshah, S. 2023. Cross-view geo-localization via learning disentangled geometric layout correspondence. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 3480–3488.
- Zhao, H.; Ren, K.; Yue, T.; Zhang, C.; and Yuan, S. 2024. TransFG: A cross-view geo-localization of satellite and UAVs imagery pipeline using transformer-based feature aggregation and gradient guidance. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–12.
- Zheng, Z.; Wei, Y.; and Yang, Y. 2020. University-1652: A multi-view multi-source benchmark for drone-based geo-localization. In *Proceedings of the 28th ACM international conference on Multimedia*, 1395–1403.
- Zhu, R.; Yin, L.; Yang, M.; Wu, F.; Yang, Y.; and Hu, W. 2023. SUES-200: A multi-height multi-scene cross-view image benchmark across drone and satellite. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(9): 4825–4839.
- Zhu, S.; Shah, M.; and Chen, C. 2022. Transgeo: Transformer is all you need for cross-view image geo-localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1162–1171.
- Zhu, S.; Yang, T.; and Chen, C. 2021. Vigor: Cross-view image geo-localization beyond one-to-one retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3640–3649.
- Zhu, Y.; Sun, B.; Lu, X.; and Jia, S. 2021. Geographic semantic network for cross-view image geo-localization. *IEEE Transactions on Geoscience and Remote Sensing*, 60: 1–15.