

# STHN: Deep Homography Estimation for UAV Thermal Geo-Localization With Satellite Imagery

Jiuhong Xiao , Graduate Student Member, IEEE, Ning Zhang , Daniel Tortei ,  
and Giuseppe Loianno , Member, IEEE

**Abstract**—Accurate geo-localization of Unmanned Aerial Vehicles (UAVs) is crucial for outdoor applications including search and rescue operations, power line inspections, and environmental monitoring. The vulnerability of Global Navigation Satellite Systems (GNSS) signals to interference and spoofing necessitates the development of additional robust localization methods for autonomous navigation. Visual Geo-localization (VG), leveraging onboard cameras and reference satellite maps, offers a promising solution for absolute localization. Specifically, Thermal Geo-localization (TG), which relies on image-based matching between thermal imagery with satellite databases, stands out by utilizing infrared cameras for effective nighttime localization. However, the efficiency and effectiveness of current TG approaches, are hindered by dense sampling on satellite maps and geometric noises in thermal query images. To overcome these challenges, we introduce STHN, a novel UAV thermal geo-localization approach that employs a coarse-to-fine deep homography estimation method. This method attains reliable thermal geo-localization within a 512-meter radius of the UAV's last known location even with a challenging 11% size ratio between thermal and satellite images, despite the presence of indistinct textures and self-similar patterns. We further show how our research significantly enhances UAV thermal geo-localization performance and robustness against geometric noises under low-visibility conditions in the wild.

**Index Terms**—Deep learning for visual perception, aerial systems: applications, localization.

## I. INTRODUCTION

THE increasing deployment of Unmanned Aerial Vehicles (UAVs) across a diverse range of applications, including agriculture [1], search and rescue operations [2], tracking [3],

Received 28 March 2024; accepted 12 August 2024. Date of publication 22 August 2024; date of current version 5 September 2024. This article was recommended for publication by Associate Editor I. Gilitschenski and Editor C. Cadena Lerma upon evaluation of the reviewers' comments. This work was supported in part by Technology Innovation Institute, in part by NSF CAREER under Award 2145277, in part by NSF CPS under Grant CNS-2121391, and in part by NYU IT High Performance Computing resources, services, and staff expertise. Giuseppe Loianno serves as consultant for the Technology Innovation Institute. This arrangement has been reviewed and approved by the New York University in accordance with its policy on objectivity in research. (Ning Zhang and Daniel Tortei contributed equally to this work.) (Corresponding author: Jiuhong Xiao.)

Jiuhong Xiao and Giuseppe Loianno are with New York University, Brooklyn, NY 11201 USA (e-mail: jx1190@nyu.edu; loiannog@nyu.edu).

Ning Zhang and Daniel Tortei are with Autonomous Robotics Research Center-Technology Innovation Institute, Abu Dhabi, UAE (e-mail: ning.zhang@tii.ae; daniel.tortei@tii.ae).

The code is made publicly available at <https://xjh19971.github.io/STHN/>. This letter has supplementary downloadable material available at <https://doi.org/10.1109/LRA.2024.3448129>, provided by the authors.

Digital Object Identifier 10.1109/LRA.2024.3448129




Fig. 1. STHN framework for UAV thermal geo-localization with satellite maps. This framework achieves robust UAV thermal localization with a challenging size ratio of 11% between thermal and satellite images.

power line inspections [4], and solar power plant inspections [5], underscores the growing importance of robust UAV localization for autonomous navigation to guarantee the effective execution of these tasks. In outdoor environments, absolute localization technology [6] is crucial as reliance on relative localization methods can cause error accumulation over time, particularly during long-time missions or in scenarios lacking loop closure detection. While Global Navigation Satellite Systems (GNSS) have become the preferred solutions, their reliability can be compromised by vulnerabilities to signal interference, jamming, and spoofing. Visual geo-localization [7], [8], [9], [10] emerges as a significant alternative solution, utilizing onboard cameras to facilitate absolute localization and navigation. This approach aligns captured RGB imagery, taken from nadir (top-down) or oblique views, with an existing reference map (such as a satellite map), enabling accurate positioning in GNSS-denied environments. However, this approach poses significant challenges in low-visibility or nighttime environments.

In response to these challenges, recent advancements in UAV thermal geo-localization [11] explore an image-based matching approach with an onboard thermal camera to match nadir-view images to satellite image crops from a database. However, this method encounters several drawbacks. Firstly, the localization accuracy is majorly influenced by the density of satellite image samples in the database. Reducing the sampling interval improves continuity between image crops and localization accuracy but increases computation time and memory usage for extensive sampling. Additionally, the approach has limited tolerance for thermal images that are not correctly north-aligned, with geometric distortions negatively impacting localization accuracy.

Addressing these limitations, this study introduces Satellite-Thermal Homography Network (STHN) framework (see Fig. 1) that leverages deep homography estimation techniques [12],

[13], [14], [15], [16], [17], [18] to directly align thermal images with satellite maps of the local region, optimizing localization in GNSS-denied scenarios. This approach adopts a two-stage coarse-to-fine strategy: 1) *Coarse alignment*, which matches small thermal images to large satellite maps within a search radius of 512 m with a challenging constant size ratio of 11%; and 2) *Refinement*, which crops and resizes the selected region and applies a second-stage estimation for enhanced accuracy.

The main contributions of this research are outlined as follows. First, we introduce for the first time a novel satellite-thermal Deep Homography Estimation (DHE) method based on an efficient coarse-to-fine approach tailored for UAV nighttime Thermal Geo-localization (TG), eliminating the dense satellite map sampling requirement of [11]. Second, we introduce the Thermal Generative Module (TGM) [11] into our DHE framework, improving the alignment between thermal and satellite images with significant scale change using limited satellite-thermal paired data. Third, we validate our approach by considering extensive and comprehensive experiments in challenging scenarios where thermal images have indistinct self-similar features on the deserts and a low overlap rate (11%) with satellite images. We demonstrate the superior performance of our method over state-of-the-art real-time DHE methods and better efficiency and accuracy over image-based matching methods. Our results also demonstrate that STHN can effectively tolerate and estimate certain geometric noises including rotation, resizing, and perspective transformation noises for thermal geo-localization. To our knowledge, this is the first deep homography estimation solution for UAV thermal geo-localization, facilitating reliable nighttime localization over long-distance outdoor flights.

## II. RELATED WORKS

*UAV Visual and Thermal Geo-localization:* UAV visual geo-localization technology has been explored by multiple works based on: 1) Template matching methods [19], [20] perform dense image alignment to optimize the image similarity measures; 2) Traditional keypoint matching methods [10], [21] extract and match the keypoints using hand-crafted detector and descriptors; and 3) Deep-learning-based matching methods [7], [8], [22], [23], [24] utilize deep neural network [25] to generate robust matching features against environmental noises. For UAV thermal localization with nadir views, [26], [27] adopt Thermal Inertial Odometry (TIO) for navigating short-distance outdoor flights. For long-distance geo-localization, [28] uses keypoint-based visible-thermal image registration, whereas [11] employs image-based matching with generative models and domain adaptation for enhanced cross-spectral geo-localization with limited training data. Despite the efficiency of keypoint-based methods, their reliance on repeatable cross-spectral local features limits their applicability. In contrast, image-based matching methods [11], [29], free from this requirement, face challenges with exhaustive searches and high memory demands, with performances that are heavily dependent on satellite database density. Our research diverges by introducing deep homography estimation for precise satellite and thermal image alignment, presenting a novel geo-localization framework that surpasses prior limitations by eliminating the necessity for repeatable local features or exhaustive searches, improving accuracy and efficiency.

*Deep Homography Estimation:* Deep homography estimation is first proposed by [13], which uses four-corner displacement as

the parametrization of homography estimation and four-corner perturbed images to train the model. [14] develops a content-aware deep homography estimation approach against the noise from the dynamic dominant foreground. [15] employs inverse compositional Lucas-Kanade algorithms for multi-modal image alignment. In [16], the authors propose LocalTrans to conduct cross-resolution homography estimation. [12] shows an iterative process to iteratively refine the homography estimation results in real-time, whereas [17] uses a focus transformer for global and local correlation to enhance estimation performance. Considering UAV localization, [30] proposes to use an unsupervised approach with photometric consistency loss for warped aerial RGB images while requiring about 65% overlap between two source images. For thermal imagery, [31] employs a multi-scale conditional GAN architecture [32] to conduct thermal-visible homography estimation. The subsequent work [33] shifts to a coarse-to-fine paradigm to further improve the estimation performance. However, the previous works commonly require a minimum overlap of 25% and, in rare instances, exactly 25%. Compared to these works, our approach adopts a coarse-to-fine paradigm but considers coarse estimation across images with major scale change for large search regions. This results in a challenging constant 11% size ratio. For refinement, our approach differs from [12], [15], [16], [33], which typically upsample aligned images. Given the small size ratio of the thermal image, a large portion of the satellite image becomes redundant and can even hinder the refinement process. Instead, we crop the selected satellite region and perform estimation without increasing image resolution to enhance efficiency.

## III. METHODOLOGY

Our STHN framework, shown in Fig. 2, has three main components: Thermal Generative Module (TGM), coarse alignment module, and refinement module.

### A. Thermal Generative Module (TGM)

We employ TGM [11] to enhance our training dataset with synthetic thermal images derived from satellite images. In the data preparation phase, we denote  $I_{OS}$  and  $I_{OT}$  as the pair of satellite and 8-bit thermal images from the original dataset, and  $I_{GS}$  as the satellite images without paired thermal images. We train TGM with the input  $I_{OS}$  and target output  $I_{OT}$  following pix2pix [34] approach. After training TGM, we generate synthetic thermal images  $I_{GT}$  using TGM and  $I_{GS}$ , and combine  $I_{OS}$  and  $I_{OT}$  to build an extended satellite-thermal dataset. We denote the quantity of actual thermal images as  $N_T$  and those generated as  $N_G$ . We restrict our sampling from the generated dataset per epoch to  $N_T$  instances to mitigate bias towards  $I_{GT}$ , given that  $N_T \ll N_G$ .

### B. Coarse-to-Fine Iterative Homography Estimation

Our coarse-to-fine strategy is divided into two stages: Coarse alignment and refinement.

1) *Coarse Alignment Stage:* We denote  $W_S$  as the width of input square satellite images  $I_S$  and  $W_T$  as that of input square 8-bit thermal images  $I_T$ . For pre-processing, we resize  $I_S$  and  $I_T$  to  $I_{RS}$  and  $I_{RT}$  at the side length of  $W_R$ , and the resize ratios of  $I_{RS}$  is  $\alpha = W_S/W_R$ . Then, we run the model

$$D_{RS \rightarrow RT} = F_H(I_{RS}, I_{RT}), \quad (1)$$




Fig. 2. STHN Framework Overview: For the data preparation phase, TGM produces synthetic thermal images from unpaired satellite images, augmenting the dataset. The deep homography estimation phase employs  $F_H$  for the Coarse Alignment Stage by predicting the displacement  $D_{RS \rightarrow RT}$  between thermal images and satellite maps. For the Refinement Stage, the framework crops and resizes the selected region  $B$ , utilizing  $F'_H$  to fine-tune the four-corner displacement prediction for enhanced accuracy.

where  $D_{RS \rightarrow RT} \in \mathbb{R}^{2 \times 4}$  is the displacement from the four corners of  $I_{RS}$  to those of  $I_{RT}$  and  $F_H$  is the homography estimation model. In other words,  $D_{RS \rightarrow RT}$  aligns  $I_{RT}$  into  $I_{RS}$ .  $F_H$  follows an iterative estimation paradigm [12], which consists of three modules: A Convolutional Neural Network (CNN) [25] feature extractor (multiple residual blocks with multi-layer CNNs and instance normalization) outputs the feature map  $\mathcal{F}_{RS}, \mathcal{F}_{RT} \in \mathbb{R}^{C \times \frac{W_R}{4} \times \frac{W_R}{4}}$  ( $C = 256$ ), a correlation module outputs correlation volumes [35]  $\mathbf{C} (\frac{W_R}{4} \times \frac{W_R}{4} \times \frac{W_R}{4} \times \frac{W_R}{4})$  and  $\mathbf{C}^{\frac{1}{2}} (\frac{W_R}{4} \times \frac{W_R}{4} \times \frac{W_R}{8} \times \frac{W_R}{8})$ , and an iterative homography estimator (multi-layer CNNs with group normalization) provides updates of displacement  $\Delta D_{RS \rightarrow RT}$ . At iteration  $k$ ,  $D_{RS \rightarrow RT}$  is updated as

$$D_{k+1, RS \rightarrow RT} = D_{k, RS \rightarrow RT} + \Delta D_{k, RS \rightarrow RT}. \quad (2)$$

Since the images are resized during pre-processing, the displacement of the coarse alignment stage on the scale of  $I_S$  is  $D'_{S \rightarrow T} = \alpha D_{RS \rightarrow RT}$ . For the loss function, we minimize the L1 distance between the predicted displacements  $D_{k, RS \rightarrow RT}$  and ground truth ones  $D_{k, RS \rightarrow RT}^{gt}$  with exponential decay as

$$\mathcal{L}_{\text{coarse}} = \sum_{k=0}^{K_1-1} \gamma^{K_1-k-1} \|D_{k, RS \rightarrow RT} - D_{k, RS \rightarrow RT}^{gt}\|_1, \quad (3)$$

where  $K_1$  is the number of updates in the coarse alignment.  $D_{RS \rightarrow RT} = D_{K_1, RS \rightarrow RT}$ . The decay factor  $\gamma$  is 0.85.

2) *Refinement Stage*: We create a bounding box  $B$  that bounds the corners of thermal images warped by  $D'_{S \rightarrow T}$ . We set  $B$  orthogonal to the image frame to ensure complete coverage of the target region, even if the coarse alignment result has rotation or perspective transformation errors. We denote  $D_{S \rightarrow B} \in \mathbb{R}^{2 \times 4}$  as the four-corner displacement from  $I_S$  to  $B$ . We crop out the

region of  $B$  to get  $I_B$  at the side length of  $W_B$  and resize it to  $I_{RB}$  at the side length of  $W_R$ . The resize ratio is  $\eta = W_B/W_R$ . The refinement process is

$$D_{RB \rightarrow RT} = F'_H(I_{RB}, I_{RT}), \quad (4)$$

where  $F'_H$  has the same structure as  $F_H$  with iterative updates (see (2)) but does not share weights and  $D_{RB \rightarrow RT} \in \mathbb{R}^{2 \times 4}$  are four-corner displacement from  $I_{RB}$  to  $I_{RT}$ . We set  $\kappa = \eta/\alpha$  and the loss function is

$$\mathcal{L}_{\text{fine}} = \sum_{k=0}^{K_2-1} \gamma^{K_2-k-1} \kappa \|D_{k, RB \rightarrow RT} - D_{k, RB \rightarrow RT}^{gt}\|_1, \quad (5)$$

where  $D_{k, RB \rightarrow RT}$  and  $D_{k, RB \rightarrow RT}^{gt}$  are predicted and ground truth displacements, and  $K_2$  is the number of updates in the refinement.  $\kappa$  maps the displacement from the scale of  $I_{RB}$  to the scale of  $I_{RS}$ , aligning with  $\mathcal{L}_{\text{coarse}}$ . The total loss function is

$$\mathcal{L} = \mathcal{L}_{\text{coarse}} + \mathcal{L}_{\text{fine}}. \quad (6)$$

The displacement of the refinement stage on the scale of  $I_S$  is  $D_{B \rightarrow T} = \eta D_{RB \rightarrow RT}$ . Combining the two stages' results, we get final displacements

$$D_{S \rightarrow T} = D_{S \rightarrow B} + D_{B \rightarrow T}. \quad (7)$$

With  $D_{S \rightarrow T}$ , we use Direct Linear Transformation (DLT) [36] to solve the homography matrix  $H \in \mathbb{R}^{3 \times 3}$ . The geo-localization center coordinate  $(x_c, y_c)$  is calculated as

$$(x_c, y_c, 1)^\top = H \times \left( \frac{W_S}{2}, \frac{W_S}{2}, 1 \right)^\top. \quad (8)$$



### C. Two-Stage Training Strategy

For training the two-stage model, we first train the coarse alignment module from scratch, and then we attach the refinement module to the end of the coarse alignment module and jointly fine-tune the two modules. We discovered that augmenting the bounding box  $B$  is crucial for effectively fine-tuning the refinement module. This requirement arises because the refinement module always tends to make no or only minor adjustments if the coarse alignment already performs well on training and validation sets. Furthermore, we observe that merely fixedly expanding the cropped boxes without random shifting and enlargement does not enhance performance. To boost the refinement module's effectiveness, we augment  $B$  by shifting the center coordinates  $(x_B, y_B)$  by  $(\Delta p_1, \Delta p_2)$  and expanding the width  $W_B$  by  $2\Delta p_3$  during training. During the evaluation phase, we consistently expand  $W_B$  by  $\Delta p_4$  to mitigate the potential offset error of the coarse alignment.

## IV. EXPERIMENTAL SETUP

### A. Dataset

For training and evaluation, our study utilizes the Boson-nighttime [11] real-world dataset which contains 10,256 train pairs, 13,011 validation pairs, and 26,568 test pairs of coupled satellite RGB and nadir-view 8-bit thermal imagery. We have expanded the dataset by augmenting the collection of satellite images without corresponding thermal images from 79,950 to 163,344 images, covering an area of 215.78 km<sup>2</sup>. This enhancement focuses on the desert and farm areas near the original dataset's sampling region, thereby incorporating a broader spectrum of geographical patterns. Additionally, the test region is then excluded from the generated data to ensure a robust evaluation of generalization performance. The thermal images in the dataset are captured between 9:00 PM and 4:00 AM, and they are aligned with an approx. spatial resolution of 1 m/px. The thermal images are cropped to  $W_T \times W_T$  pixels (px), where  $W_T = 512$ . The satellite images<sup>1</sup> are cropped to  $W_S \times W_S$ . Fig. 3 shows the ground truth overlap between thermal images and satellite images with different  $W_S$ . For  $W_S = 512/1024/1536$ , the size ratios between thermal images and satellite images are 100%, 25%, and 11%.

### B. Metrics

We deploy two accuracy metrics in our evaluation: Mean Average Corner Error (MACE) and Center Error (CE). MACE, extensively adopted in [12], [13], [17], measures the mean value of the average distances between the four corners of estimated and ground truth image alignments. Conversely, CE measures the mean value of the distances between the center points of predicted thermal image displacements and ground truth ones, thereby measuring geo-localization accuracy.

In our experimental analysis, the maximum spatial distance between the center points of input thermal and satellite images  $D_C$  emerges as a critical factor influencing estimation performance. Intuitively, a larger  $D_C$  implies a greater translation from the center required for the four-corner displacement, which in turn becomes more challenging to predict accurately. To validate




Fig. 3. Example images of Boson-nighttime dataset. The 1st row and 2nd row are input satellite and thermal images. The 3rd–5th rows are the ground truth overlap between satellite and thermal images with different  $W_S$ .

the robustness of our method, we cautiously ablate results across a spectrum of  $D_C$ , demonstrating our approach's capability under varying degrees of challenging translations.

### C. Implementation Details

For pre-processing, the resize side length  $W_R$  is 256 px. The training iteration numbers of the coarse alignment and refinement modules are 200000 with a batch size of 16. The AdamW optimizer [37] is employed for model training, utilizing a linear learning rate decay scheduler with warmup with the peak learning rate at  $1e-4$ . The numbers of iterative updates  $K_1$  and  $K_2$  are both set to 6. Depending on the setting, the correlation module's level is 2 (for  $W_S = 512$ ) or 4 (for  $W_S = 1024, 1536$ ) with a search radius of 4. For bounding box augmentation,  $\Delta p_1, \Delta p_2$  is set to vary between  $(-64, 64)$ ,  $\Delta p_3$  is set within  $[0, 64)$ , and  $\Delta p_4$  is 64 by parameter tuning. For geometric noises, we extend the coverage of thermal images, apply corresponding data augmentations, and center crop the thermal images to avoid black padding on their boundary. Our models are developed using PyTorch. The inference speed is measured with one NVIDIA RTX-2080-Ti GPU.

## V. RESULTS

In Sections V-A and V-B, we assume that thermal images are aligned to the north, facilitated by an onboard compass and a gimbaled thermal camera. Subsequently, in Section V-C, we broaden our analysis for geometric noises.

<sup>1</sup>Bing RGB satellite imagery is sourced from Maxar: <https://www.bing.com/maps/aerial>

TABLE I  
COMPARISON OF TEST MACE (M) BETWEEN DIFFERENT HOMOGRAPHY ESTIMATION METHODS ACROSS DIFFERENT  $D_C$

| Methods                                      | $D_C = 50$ m | $D_C = 64$ m | $D_C = 128$ m | $D_C = 256$ m | $D_C = 512$ m | Failure Rate |
|----------------------------------------------|--------------|--------------|---------------|---------------|---------------|--------------|
| <i>Traditional Keypoint Matching Methods</i> |              |              |               |               |               |              |
| Identity                                     | 35.63        | 39.08        | 85.63         | 170.94        | 334.68        | -            |
| SIFT [38] + RANSAC [39]                      | 442.20       | 654.77       | 547.29        | 529.63        | 1650.46       | 99.6%        |
| SIFT [38] + MAGSAC++ [40]                    | 512.60       | 438.54       | 529.46        | 561.64        | 693.03        | 99.7%        |
| ORB [41] + RANSAC [39]                       | 720.80       | 733.69       | 733.94        | 4614.84       | 975.83        | 82.6%        |
| ORB [41] + MAGSAC++ [40]                     | 784.12       | 558.51       | 564.63        | 524.99        | 573.72        | 82.9%        |
| BRISK [42] + RANSAC [39]                     | 503.25       | 771.21       | 665.94        | 974.80        | 591.17        | 95.5%        |
| BRISK [42] + MAGSAC++ [40]                   | 536.52       | 487.95       | 722.76        | 1948.99       | 568.24        | 95.6%        |
| <i>Learned Keypoint Matching Methods</i>     |              |              |               |               |               |              |
| R2D2 [43] + RANSAC [39]                      | 994.31       | 1160.73      | 1160.21       | 2400.31       | 2902.81       | 88.9%        |
| LoFTR [43] + RANSAC [39]                     | 1123.74      | 1697.33      | 1317.69       | 1269.71       | 2564.65       | 0%           |
| <i>Deep Homography Estimation Methods</i>    |              |              |               |               |               |              |
| DHN [13]                                     | 16.78        | 20.43        | 77.68         | 197.27        | 457.23        | 0%           |
| LocalTrans [16]                              | 33.31        | 37.29        | 86.04         | 166.52        | 338.21        | 0%           |
| IHN [12]                                     | 5.91         | 7.81         | 51.74         | 190.93        | 367.24        | 0%           |
| Ours ( $W_S = 512$ )                         | <b>4.24</b>  | <b>4.93</b>  | 14.97         | 142.71        | 347.50        | 0%           |
| Ours ( $W_S = 1024$ )                        | 4.92         | 5.31         | <b>6.03</b>   | <b>9.22</b>   | 86.74         | 0%           |
| Ours ( $W_S = 1536$ )                        | 6.50         | 7.04         | 7.27          | 16.78         | 16.42         | 0%           |
| Ours ( $W_S = 1536$ + two stages)            | 7.51         | 7.20         | 7.51          | 14.99         | <b>12.70</b>  | 0%           |

"Identity" indicates the error if no homography estimation is applied. If not specified, the methods are evaluated with  $W_S = 512$ . The bold values represent the best results.

#### A. Comparison With Baselines

In the results detailed in Table I, we initiate the analysis by evaluating the efficacy of traditional keypoint matching methods, such as SIFT [38], ORB [41], and BRISK [42], integrated with outlier rejection methods like RANSAC [39] and MAGSAC++ [40]. We also evaluate learned keypoint methods including R2D2 [43] trained on our dataset and LoFTR [46] with pretrained weights. These methods demonstrate a significantly high MACE alongside substantial failure rates (calculated by instances where the number of matching keypoints  $\leq 10$ ). This underlines the challenges of keypoint matching inherent in complex satellite-thermal alignment.

Subsequently, our analysis compares our methods with various deep homography estimation frameworks, including DHN [13], LocalTrans [16], and IHN [12] (state-of-the-art method in real-time applications). These baselines with one-stage models are trained on the Boson-nighttime dataset. We report the baseline results considering  $W_S = 512$  as representative results since other  $W_S$  show similar trends in our analysis. The results show the superior performance of our approach for satellite-thermal alignment and geo-localization. A notable observation from the data is the different performance preferences across varying  $D_C$  distances: for  $D_C = 50$  m and  $D_C = 64$  m, the optimal  $W_S$  is 512, while for mid-range distances of  $D_C = 128$  m and  $D_C = 256$  m, using  $W_S = 1024$  leads to the best results. Additionally, for the longest distance of  $D_C = 51$  m, our novel two-stage method with  $W_S = 1536$  emerges as the most effective strategy. The findings indicate that for cases where  $D_C \leq 256$  m, employing our one-stage method combined with a carefully chosen  $W_S$  emerges as the most effective strategy. Further explanation of the correlation between  $W_S$  and  $D_C$  is in Section V-B.

We find that our two-stage method fails to enhance performance for distances  $D_C = 50$  m, 64 m, and 128 m, instead leading to a decline in accuracy. Upon examining the visualized outcomes, we observe that for smaller distances ( $D_C \leq 128$  m), the initial coarse alignment is sufficiently accurate, making the

TABLE II  
COMPARISON BETWEEN DIFFERENT IMAGE-BASED MATCHING METHODS AND OUR ESTIMATION METHOD WHEN  $D_C = 512$  M

| Methods                                   | Test CE (m)  | Latency (ms) |
|-------------------------------------------|--------------|--------------|
| <i>Image-based Matching Methods</i>       |              |              |
| AnyLoc-VLAD-DINOv2 [29]                   | 258.21       | 352404.03    |
| STGL-NetVLAD-ResNet50 [11], [44]          | 89.31        | 7180.0       |
| STGL-GeM-ResNet50 [11], [45]              | 13.52        | 4918.9       |
| <i>Deep Homography Estimation Methods</i> |              |              |
| Ours ( $W_S = 1536$ )                     | 15.90        | <b>35.2</b>  |
| Ours ( $W_S = 1536$ + two stages)         | <b>12.12</b> | 63.9         |

The bold values represent the best results.

refinement module's excessive iterative updates introduce noise into the final predictions, thereby degrading performances. Nevertheless, our two-stage approach maintains an overall MACE of less than 15 m across all considered  $D_C$ , establishing robust baselines for this task. Notably, for achieving precise geo-localization at  $D_C = 512$  m, this two-stage strategy demonstrates the best performance, underscoring its effectiveness for large-scale search regions.

We also compare with image-based solutions (AnyLoc [29] and STGL [11]) on accuracy and latency aspects in Table II. The latency of image-based matching methods is calculated by  $t_e \times (N_S + 1) + t_m$ , where  $t_e$  is feature extraction time per image, and  $N_S = 841$  is the number of database images centered within a  $1024 \times 1024$  area (while the complete images cover a  $1536 \times 1536$  area) with a sampling stride of 35 px following [11], and  $t_m$  is the matching time per query. For AnyLoc, we directly apply the original DINOv2 [47] weights and fit the VLAD [48] parameters using our training data. We observe a significant performance decline in AnyLoc, likely due to the domain gap between satellite and thermal imagery. STGL with GeM yields high accuracy but still suffers from high latency. Our method exhibits significant enhancements in both accuracy and latency compared to these existing image-based matching techniques. Notably, our one-stage and two-stage methods achieve latency




Fig. 4. Effectiveness of TGM in deep homography estimation across different  $D_C$  when  $W_S = 512$ . Validation MACE (Val MACE) is plotted on a log scale.




Fig. 5. Coarse alignment under large-scale ( $W_S = 1536$ ), median-scale ( $W_S = 1024$ ), small-scale ( $W_S = 512$ ) satellite images with TGM.

reductions to just 7.2% and 13.0% of the latency of STGL-GeM-ResNet50.

### B. Ablation Study

In this study (Figs. 4, 5, and 6), we focus on the following questions

- How does the incorporation of TGM affect the accuracy of homography estimation across varying  $D_C$ ?
- Is the coarse alignment effective in achieving satisfactory localization accuracy for large  $D_C$ ?
- Is the bounding box augmentation effective for fine-tuning the refinement module?

1) *Effectiveness of TGM*: Fig. 4 demonstrates the effectiveness of TGM in improving deep homography estimation over different spatial distances between centers ( $D_C$ ) on the validation set. It showcases TGM's ability to enhance estimation accuracy by generating synthetic thermal images for satellite imagery that lacks paired thermal data. This consistent enhancement in image-based matching [11] and deep homography estimation for




Fig. 6. Qualitative comparison between finetuning the refinement module without bbox aug, with only bbox exp, and with bbox aug with  $W_S = 1536$ . Green boxes are the ground truth, blue boxes are the bounding boxes from coarse alignment, and red boxes are final predictions after refinement.

satellite-thermal matching suggests TGM's potential applicability in additional computer vision tasks that do not have direct thermal imaging counterparts.

2) *Coarse Alignment*: Fig. 5 illustrates the correlation between validation MACE and  $W_S$  across various  $D_C$ . The figure shows that as  $W_S$  increases, the validation MACE for smaller translation distances ( $D_C = 50$  m and  $64$  m) slightly increases, suggesting a deterioration in alignment accuracy. In contrast, for larger translation distances ( $D_C = 128$  m,  $256$  m,  $512$  m), the validation MACE decreases, indicating improved alignment accuracy. The intuition is that an increase in  $W_S$ , without a corresponding adjustment in  $W_R$ , leads to a higher pixel-per-meter (ppm) ratio after image resizing. This increment in ppm ratio can negatively affect the alignment accuracy. Conversely, a larger  $W_S$  enhances alignment accuracy for greater translation ( $D_C$ ), especially for  $W_S = 1536$  and  $D_C = 512$  m. In these cases, a larger  $W_S$  ensures the full coverage of the thermal image, which is crucial for accurately calculating correlation volumes  $C$ .

3) *Effectiveness of Bounding Box Augmentation*: We present a qualitative comparison in Fig. 6 to demonstrate the impact of fine-tuning with and without bounding box augmentation (bbox aug). Given that bounding box augmentation requires an expansion of the bounding box (bbox exp) during the evaluation phase, we also include results featuring solely bbox exp without bbox aug to ablate the effects. The findings illustrate that in the absence of augmentation, the refinement module tends to make only minimal adjustments when not trained with bbox exp. On the other hand, if we train the refinement module with only bbox exp, it always tends to reduce the size of the predicted box towards the center, rather than correctly repositioning it. However, the incorporation of augmentation addresses these limitations by augmenting the width and the center coordinates of the region.

### C. Robustness Evaluation and Visualization

Ideally, the UAV onboard compass and gimbal camera would supply precise data, enabling the accurate alignment of images to the north. However, it is crucial for our algorithm to






Fig. 7. Visualization results with geometric noises for our one-stage and two-stage methods with  $W_S = 1536$  and  $D_C = 512$  m. Green boxes are the ground truth, blue boxes are the bounding boxes from coarse alignment for our two-stage method, and red boxes are the final predictions of one-stage and two-stage methods. The 1st–2nd columns show rotation noises, the 3rd–4th columns show resizing noises, and the 5th–8th columns show perspective transformation noises.

TABLE III  
ROBUSTNESS EVALUATION WITH GEOMETRIC NOISES, INCLUDING ROTATION, RESIZING, AND PERSPECTIVE TRANSFORMATION NOISES WHEN  $D_C = 512$  M AND  $W_S = 1536$

|                                          | Ours (one stage) |             | Ours (two stages) |             |
|------------------------------------------|------------------|-------------|-------------------|-------------|
|                                          | Test MACE (m)    | Test CE (m) | Test MACE (m)     | Test CE (m) |
| Baseline                                 | 16.42            | 15.90       | 12.70             | 12.12       |
| <i>Rotation Noises</i>                   |                  |             |                   |             |
| 5°                                       | 31.82            | 27.00       | 14.48             | 11.42       |
| 10°                                      | 33.78            | 28.40       | 14.55             | 11.90       |
| 30°                                      | 64.63            | 56.98       | 33.48             | 28.93       |
| <i>Resizing Noises</i>                   |                  |             |                   |             |
| 0.1                                      | 32.00            | 27.38       | 13.03             | 11.20       |
| 0.2                                      | 36.79            | 31.88       | 15.16             | 13.23       |
| 0.3                                      | 46.22            | 40.50       | 21.68             | 19.36       |
| <i>Perspective Transformation Noises</i> |                  |             |                   |             |
| 8 px                                     | 30.53            | 26.41       | 13.70             | 10.94       |
| 16 px                                    | 34.07            | 27.22       | 16.24             | 11.73       |
| 32 px                                    | 44.03            | 30.35       | 20.90             | 14.11       |

“Baseline” is our method with only translation.

demonstrate tolerance towards certain rotation and perspective transformation inaccuracies during active flights. Additionally, understanding how our algorithm performs when there is a change in flight altitude—which results in a change of the thermal image’s coverage area, denoted as resizing noise—is essential. To assess the algorithm’s robustness under these conditions, we perform experiments that introduce specific rotation, resizing, and perspective transformation noises. For rotation disturbances, the thermal images undergo random rotations up to 5°, 10°, or 30°. For resizing disturbances, the images are randomly scaled by a factor of  $1 + \Delta r$ , with  $\Delta r$  varying within either  $\pm 0.1$ ,  $\pm 0.2$ , or  $\pm 0.3$ . For perspective transformation, we randomly adjust the four corners of  $512 \times 512$  thermal images up to 8 px, 16 px, or 32 px.

In Table III, we evaluate the robustness of our one-stage and two-stage strategies against a variety of geometric noise conditions with  $D_C = 512$  m and  $W_S = 1536$ . The analysis indicates a significant decrease in performance for the one-stage method under these conditions, in contrast to the two-stage strategy, which demonstrates a notable robustness against geometric perturbations. Specifically, the two-stage strategy effectively maintains test MACE below 22 m and test CE below 20 m in most scenarios, with notable exceptions being in instances of 30° rotation noise. While incremental perspective transformations and resizing have minimal impact on accuracy, large rotation

noise can significantly degrade performance. This suggests the tolerance of our strategies to different types of geometric noise. Overall, the results validate our method’s robustness and its ability to estimate these disturbances, underscoring the two-stage strategy’s superior effectiveness and reliability in mitigating the negative effects of these disturbances. Fig. 7 further illustrates this point by showcasing visual comparisons between the failure instances of the one-stage method and the success cases of the two-stage method, demonstrating the latter’s improved robustness.

## VI. CONCLUSION

This letter presents a novel deep homography estimation approach for UAV thermal geo-localization tasks. We validate the capability of STHN to precisely align thermal images, captured by UAV onboard sensors, with large-scale satellite maps, achieving successful alignment even with a size ratio of 11%. Additionally, we showcase STHN’s superior performances in terms of speed and accuracy with respect to several state-of-the-art approaches as well as its resilience to geometric distortions, which significantly enhances the reliability of geo-localization outcomes.

Our future endeavors will aim to develop a hierarchical geo-localization framework. This framework will integrate deep homography estimation for local matching with image-based matching techniques for broad-scale global matching, thereby building up universal geo-localization solutions.

## REFERENCES

- [1] D. C. Tsouros, S. Bibi, and P. G. Sarigiannidis, “A review on UAV-based applications for precision agriculture,” *Information*, vol. 10, no. 11, 2019, Art. no. 349.
- [2] M. Pólka, S. Ptak, and Łukasz Kuziora, “The use of UAV’s for search and rescue operations,” in *Procedia Eng.*, vol. 192, pp. 748–752, 2017.
- [3] A. Saviolo, P. Rao, V. Radhakrishnan, J. Xiao, and G. Loianno, “Unifying foundation models with quadrotor control for visual tracking beyond object categories,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2024, pp. 7389–7396.
- [4] J. Xing, G. Cioffi, J. Hidalgo-Carrió, and D. Scaramuzza, “Autonomous power line inspection with drones via perception-aware MPC,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2023, pp. 1086–1093.
- [5] L. Morando, C. T. Recchiuto, J. Calla, P. Scuteri, and A. Sgorbissa, “Thermal and visual tracking of photovoltaic plants for autonomous UAV inspection,” in *Drones*, vol. 6, no. 11, 2022, Art. no. 347.
- [6] A. Couturier and M. A. Akhloufi, “A review on absolute visual localization for UAV,” *Robot. Auton. Syst.*, vol. 135, 2021, Art. no. 103666.

- [7] Y. He et al., "Foundloc: Vision-based onboard aerial localization in the wild," 2023, *arXiv:2310.16299*.
- [8] A. T. Fragoso, C. T. Lee, A. S. McCoy, and S.-J. Chung, "A seasonally invariant deep transform for visual terrain-relative navigation," *Sci. Robot.*, vol. 6, no. 55, 2021, Art. no. eabf3320.
- [9] B. Patel, T. D. Barfoot, and A. P. Schoellig, "Visual localization with Google Earth images for robust global pose estimation of UAVs," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2020, pp. 6491–6497.
- [10] M. Shan, F. Wang, F. Lin, Z. Gao, Y. Z. Tang, and B. M. Chen, "Google Map aided visual navigation for UAVs in GPS-denied environment," in *Proc. IEEE Int. Conf. Robot. Biomimetics*, 2015, pp. 114–119.
- [11] J. Xiao, D. Tortei, E. Roura, and G. Loianno, "Long-range UAV thermal geo-localization with satellite imagery," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2023, pp. 5820–5827.
- [12] S.-Y. Cao, J. Hu, Z. Sheng, and H.-L. Shen, "Iterative deep homography estimation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 1879–1888.
- [13] D. DeTone, T. Malisiewicz, and A. Rabinovich, "Deep image homography estimation," 2016, *arXiv:1606.03798*.
- [14] J. Zhang et al., "Content-aware unsupervised deep homography estimation," in *Proc. 16th Eur. Conf. Comput. Vis.*, 2020, pp. 653–669.
- [15] Y. Zhao, X. Huang, and Z. Zhang, "Deep Lucas-Kanade homography for multimodal image alignment," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 15950–15959.
- [16] R. Shao, G. Wu, Y. Zhou, Y. Fu, L. Fang, and Y. Liu, "Localtrans: A multiscale local transformer network for cross-resolution homography estimation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 14890–14899.
- [17] S.-Y. Cao et al., "Recurrent homography estimation using homography-guided image warping and focus transformer," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2023, pp. 9833–9842.
- [18] H. Le, F. Liu, S. Zhang, and A. Agarwala, "Deep homography estimation for dynamic scenes," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 7652–7661.
- [19] G. J. V. Dalen, D. P. Magree, and E. N. Johnson, "Absolute localization using image alignment and particle filtering," in *Proc. AIAA Guid. Navigation Control Conf.*, 2016, Art. no. 0647.
- [20] A. Yol, B. Delabarre, A. Dame, J.-E. Dartois, and E. Marchand, "Vision-based absolute localization for unmanned aerial vehicles," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2014, pp. 3429–3434.
- [21] M. Mantelli et al., "A novel measurement model based on abbRIEF for global localization of a UAV over satellite images," *Robot. Auton. Syst.*, vol. 112, pp. 304–319, 2019.
- [22] A. Shetty and G. X. Gao, "UAV pose estimation using cross-view geolocalization with satellite imagery," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2019, pp. 1827–1833.
- [23] H. Goforth and S. Lucey, "GPS-denied UAV localization using pre-existing satellite imagery," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2019, pp. 2974–2980.
- [24] M. Bianchi and T. D. Barfoot, "UAV localization using autoencoded satellite images," *IEEE Robot. Automat. Lett.*, vol. 6, no. 2, pp. 1761–1768, Apr. 2021.
- [25] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 27553, pp. 436–444, 2015.
- [26] J. Delaune, R. Hewitt, L. Lytle, C. Sorice, R. Thakker, and L. Matthies, "Thermal-inertial odometry for autonomous flight throughout the night," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2019, pp. 1122–1128.
- [27] V. Polizzi, R. Hewitt, J. Hidalgo-Carri , J. Delaune, and D. Scaramuzza, "Data-efficient collaborative decentralized thermal-inertial odometry," *IEEE Robot. Automat. Lett.*, vol. 7, no. 4, pp. 10681–10688, Oct. 2022.
- [28] F. Achermann et al., "MultiPoint: Cross-spectral registration of thermal and optical aerial imagery," in *Proc. Conf. Robot Learn.*, Nov. 2021, pp. 1746–1760.
- [29] N. Keetha et al., "AnyLoc: Towards universal visual place recognition," *IEEE Robot. Automat. Lett.*, vol. 9, no. 2, pp. 1286–1293, Feb. 2024.
- [30] T. Nguyen, S. W. Chen, S. S. Shivakumar, C. J. Taylor, and V. Kumar, "Unsupervised deep homography: A fast and robust homography estimation model," *IEEE Robot. Automat. Lett.*, vol. 3, no. 3, pp. 2346–2353, Mar. 2018.
- [31] Y. Luo, X. Wang, Y. Wu, and C. Shu, "Infrared and visible image homography estimation using multiscale generative adversarial network," *Electronics*, vol. 12, no. 4, 2023, Art. no. 788.
- [32] M. Mirza and S. Osindero, "Conditional generative adversarial nets," 2014, *arXiv:1411.1784*.
- [33] X. Wang et al., "Coarse-to-fine homography estimation for infrared and visible images," *Electronics*, vol. 12, no. 21, 2023, Art. no. 4441.
- [34] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 1125–1134.
- [35] Z. Teed and J. Deng, "RAFT: Recurrent all-pairs field transforms for optical flow," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 402–419.
- [36] Y. I. Abdel-Aziz, H. M. Karara, and M. Hauck, "Direct linear transformation from comparator coordinates into object space coordinates in close-range photogrammetry," *Photogrammetric Eng. Remote Sens.*, vol. 81, no. 2, pp. 103–107, 2015.
- [37] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *Proc. Int. Conf. Learn. Representations*, 2019.
- [38] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *Int. J. Comput. Vis.*, vol. 60, no. 2, pp. 91–110, 2004.
- [39] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, pp. 381–395, Jun. 1981.
- [40] D. Bar th, J. Noskova, M. Ivashechkin, and J. Matas, "MAGSAC++, a fast, reliable and accurate robust estimator," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 1301–1309.
- [41] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "ORB: An efficient alternative to sift or surf," in *Proc. Int. Conf. Comput. Vis.*, 2011, pp. 2564–2571.
- [42] S. Leutenegger, M. Chli, and R. Y. Siegwart, "BRISK: Binary robust invariant scalable keypoints," in *Proc. Int. Conf. Comput. Vis.*, 2011, pp. 2548–2555.
- [43] J. Revaud, P. Weinzaepfel, C. R. de Souza, and M. Humenberger, "R2D2: Repeatable and reliable detector and descriptor," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 12405–12415.
- [44] R. Arandjelovi , P. Gronat, A. Torii, T. Pajdla, and J. Sivic, "NetVLAD: CNN architecture for weakly supervised place recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 6, pp. 1437–1451, Jun. 2018.
- [45] F. Radenovi , G. Tolias, and O. Chum, "Fine-tuning CNN image retrieval with no human annotation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 7, pp. 1655–1668, Jul. 2019.
- [46] J. Sun, Z. Shen, Y. Wang, H. Bao, and X. Zhou, "LoFTR: Detector-free local feature matching with transformers," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 8918–8927.
- [47] M. Oquab et al., "DINOv2: Learning robust visual features without supervision," *Trans. Mach. Learn. Res.*, 2024. [Online]. Available: <https://openreview.net/forum?id=a68SUt6zFt>
- [48] H. J gou, M. Douze, C. Schmid, and P. P rez, "Aggregating local descriptors into a compact image representation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2010, pp. 3304–3311.