

# LST-SLAM: A Stereo Thermal SLAM System for Kilometer-Scale Dynamic Environments

Zeyu Jiang<sup>1</sup>, Kuan Xu<sup>2</sup>, Changhao Chen<sup>1\*</sup>

**Abstract**—Thermal cameras offer strong potential for robot perception under challenging illumination and weather conditions. However, thermal Simultaneous Localization and Mapping (SLAM) remains difficult due to unreliable feature extraction, unstable motion tracking, and inconsistent global pose and map construction, particularly in dynamic large-scale outdoor environments. To address these challenges, we propose LST-SLAM, a novel large-scale stereo thermal SLAM system that achieves robust performance in complex, dynamic scenes. Our approach combines self-supervised thermal feature learning, stereo dual-level motion tracking, and geometric pose optimization. We also introduce a semantic-geometric hybrid constraint that suppresses potentially dynamic features lacking strong inter-frame geometric consistency. Furthermore, we develop an online incremental bag-of-words model for loop closure detection, coupled with global pose optimization to mitigate accumulated drift. Extensive experiments on kilometer-scale dynamic thermal datasets show that LST-SLAM significantly outperforms recent representative SLAM systems, including AirSLAM and DROID-SLAM, in both robustness and accuracy.

## I. INTRODUCTION

Visual Simultaneous Localization and Mapping (SLAM) enables autonomous robots to estimate their motion in real time while simultaneously building a map of their surroundings using cameras. Owing to its balance between cost and functionality, visual SLAM has been widely adopted in various robotic applications, such as service and delivery robots [1]. However, RGB camera-based visual SLAM often suffers from severe performance degradation in challenging illumination conditions, such as low-light environments or sudden lighting changes.

Thermal cameras provide a promising alternative. Unlike RGB cameras, which rely on reflected visible light, thermal cameras capture infrared radiation emitted by objects, allowing robust imaging in poor lighting conditions and even in complete darkness. They also demonstrate greater resilience in environments affected by fog, dust, or smoke [2]. Consequently, thermal camera-based SLAM systems [3], [4] have demonstrated strong potential under illumination-challenging conditions. Nonetheless, most existing thermal

<sup>1</sup>Zeyu Jiang and Changhao Chen are with PEAK-Lab, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, 511453, China zjiang122@connect.hkust-gz.edu.cn, changhaochen@hkust-gz.edu.cn

<sup>2</sup>Kuan Xu is with the School of Electrical and Electronic Engineering, Nanyang Technological University, 50 Nanyang Avenue, Singapore 639798 kuan.xu@ntu.edu.sg

This work was supported by National Natural Science Foundation of China (NSFC) under the Grant Number 62573370 and Key Area Project of Education Department of Guangdong Province (No. 2025ZDZX3051).

\* Corresponding Author.




Fig. 1: LST-SLAM enables robust localization and mapping in kilometer-scale dynamic thermal scenes. The system leverages self-supervised thermal features and a dual-level tracking strategy, while an incremental online BoW supports loop closure detection and global pose optimization.

SLAM systems rely on feature extraction and tracking methods originally designed for RGB images, despite the fact that thermal imaging is physically different and inherently suffers from low contrast, weak texture, and high noise. Moreover, non-uniformity correction (NUC) introduces inter-frame inconsistencies [5], making it difficult to extract sufficiently stable features for long-term motion tracking.

To address these limitations, several works have fused thermal cameras with additional sensors such as LiDAR [6], RGB cameras [7], or IMUs [8]. While effective, these approaches increase hardware costs and require substantial calibration and synchronization effort, reducing their practicality and leaving the fundamental challenge of stable thermal-only motion tracking unresolved. Furthermore, realizing robust thermal SLAM in large-scale outdoor environments remains an underexplored area. Outdoor scenes are typically unstructured, offering fewer stable features than indoor environments, making visual tracking prone to frequent failures. The prevalence of dynamic objects further exacerbates visual aliasing, reducing the robustness of feature tracking. During back-end pose optimization, the outliers introduced by dynamic objects further degrade system accuracy [4]. Together, these challenges significantly compromise the robustness and accuracy of existing thermal SLAM systems in large-scale scenarios.

In this work, we propose LST-SLAM, a novel Large-scale

Stereo Thermal SLAM system designed to achieve reliable performance in kilometer-scale, illumination-challenging, and dynamic outdoor environments. LST-SLAM integrates thermal feature self-learning, dual-level motion tracking, and geometric pose optimization. To enable stable thermal feature extraction, we introduce a thermal point network with a thermal homography training loss that adapts geometric knowledge learned in the RGB domain to the thermal domain. To improve motion tracking robustness, we design a stereo dual-level feature-tracking strategy that jointly optimizes low-level photometric and high-level descriptor errors. In addition, we leverage semantic and geometric cues to perform dynamic feature filtering, mitigating the impact of moving objects commonly present in real-world traffic scenes. Finally, we construct an incremental bag-of-words (BoW) of thermal features for loop closure detection, coupled with global pose optimization to reduce accumulated drift. In real-world kilometer-scale experiments, LST-SLAM achieves 75.8% and 66.8% lower localization errors than representative systems, i.e., AirSLAM (2025) and DROID-SLAM (2021), respectively.

In summary, our main contributions are as follows:

- We propose LST-SLAM, a novel thermal SLAM system tailored for large-scale, dynamic, and illumination-challenging outdoor scenarios, featuring thermal feature self-learning, dual-level motion tracking, incremental thermal BoW construction, and geometric pose optimization.
- We introduce a thermal point network with a thermal homography training loss that adapts geometric knowledge from the RGB domain to the thermal domain. Together with a stereo dual-level tracking strategy and dynamic feature filtering, our system transforms learned thermal features into reliable motion tracking.
- We conduct extensive experiments in highly challenging kilometer-scale outdoor environments, demonstrating that LST-SLAM achieves superior feature tracking stability and localization accuracy compared to several representative SLAM systems.

## II. RELATED WORKS

**Thermal Feature Extractor** Mouats *et al.*[9] established a benchmark thermal dataset to evaluate the performance of handcrafted feature extractors. However, traditional detectors and descriptors [10], [11], [12], initially designed for visible-spectrum images, are difficult to adapt to thermal data since they rely heavily on image gradient information. To address these limitations, deep learning approaches have been explored. Zhao *et al.*[8] proposed ThermalPoint, a supervised method trained on authentic thermal images, which mitigates weak texture and illumination variations while improving noise robustness. Lu *et al.*[13] introduced a complete learning pipeline for robust thermal feature detection and description. Deshpande *et al.*[14] designed a triplet-based Siamese CNN to extract features from arbitrary thermal images. More recently, Tuzcuoglu *et al.*[15] developed a cross-modal, cross-view approach for local feature matching

between thermal infrared and visible images. Such cross-modal strategies offer valuable insights for advancing thermal feature learning.

**Thermal SLAM** Compared to RGB SLAM, thermal SLAM has received relatively limited attention. Most prior efforts have relied on multi-sensor fusion strategies, combining thermal cameras with LiDAR [6], [16], IMU [8], [17], [18], or RGB cameras [19], [7] to leverage complementary sensing modalities for improved robustness in extreme conditions. Early work by Vidas *et al.*[2] presented the first SLAM system using a single thermal sensor, reporting large reprojection errors that highlighted challenges in thermal feature matching. More recently, Wu *et al.*[4] proposed a monocular thermal SLAM method that incorporates dynamic object segmentation to mitigate feature degradation in dynamic environments. Xu *et al.*[20] introduced an end-to-end monocular thermal SLAM system that achieves accurate localization and dense 3D mapping in nighttime conditions. Gupta *et al.*[3] investigated feature matching and loop closure under diurnal temperature variations, proposing a long-term thermal SLAM framework. Wu *et al.*[21] further integrated Neural Radiance Fields (NeRF) into thermal SLAM for thermal-based 3D reconstruction. These methods have improved RGB-based SLAM systems in specific modules, but they have not been designed for thermal features across the entire system. Therefore, previous research struggles to be applied to kilometer-scale real-world thermal scenarios.

## III. METHODOLOGY

As illustrated in Fig. 2, LST-SLAM combines a learning-based thermal feature detection module with dynamic object filtering, an incremental online BoW framework for loop closure, and a pose optimization backend. Pre-processed thermal images (Section III-A) are first fed into the self-supervised thermal point network described in Section III-B. Robust stereo static thermal feature keypoints are then obtained via the dynamic object filtering model (Section III-C). These keypoints are used in stereo dual-level motion tracking (Section III-D) and bundle adjustment (BA) (Section III-E) to refine both camera poses and the reconstructed map. Finally, during loop closure (Section III-F), the incremental online BoW scheme enables generalized loop detection and global pose optimization.

### A. Thermal Image Preprocess

Raw infrared images are typically 16-bit, low-contrast, and noisy, making them unsuitable for direct feature extraction. To address this, we apply a preprocessing pipeline to enhance image quality before subsequent processing. Specifically, each raw 16-bit thermal frame is first normalized to 8-bit using percentile stretching, with intensity bounds defined by the 1% and 99%. This suppresses outliers while preserving the main scene content. To reduce flickering caused by frame-to-frame variations, the percentile bounds are smoothed over time using an exponential moving average (EMA):

$$l_t = \alpha l_{t-1} + (1 - \alpha)\hat{l}_t, \quad h_t = \alpha h_{t-1} + (1 - \alpha)\hat{h}_t \quad (1)$$

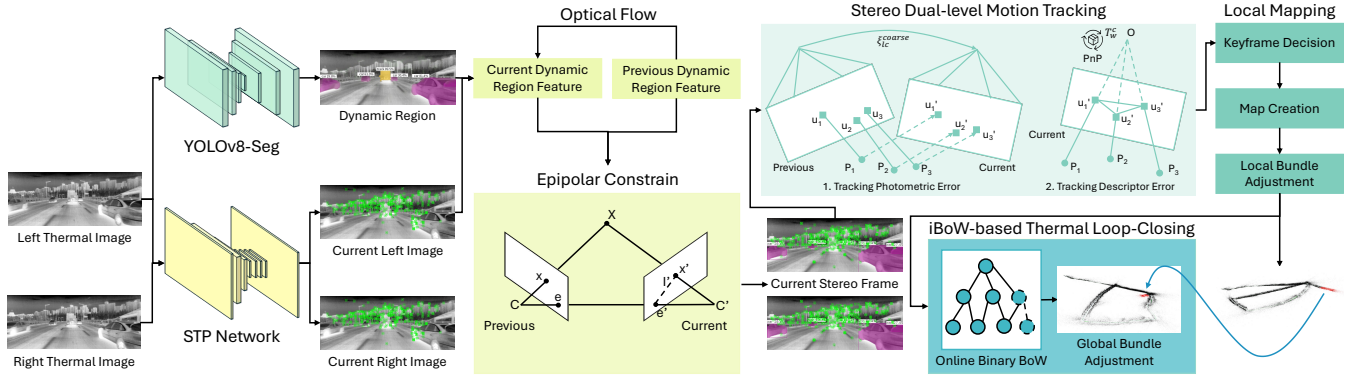


Fig. 2: LST-SLAM system pipeline. Stereo inputs are processed with self-supervised feature point and dynamic filtering networks to obtain robust thermal features, which drive tracking, mapping, and local optimization. Binarized descriptors build an incremental BoW (iBoW) for loop closure and global pose optimization.

where  $\hat{l}_t$  and  $\hat{h}_t$  denote the raw percentile estimates, and  $\alpha \in [0, 1)$  is the smoothing factor. We observed negligible performance differences in the range  $\alpha \in [0.7, 0.9]$ . The stretched image is then linearly mapped to the range  $[0, 255]$ .

Finally, contrast-limited adaptive histogram equalization (CLAHE) is applied to enhance local contrast, with the clip limit and tile size chosen to balance detail enhancement against noise amplification. This preprocessing pipeline produces temporally stable, high-contrast thermal images that are well-suited for both visualization and downstream computer vision tasks.

### B. Self-Supervised Thermal Feature Learning

In feature-based visual SLAM, the quality of feature descriptors is critical for data association. High-quality descriptors enable reliable feature matching, which directly impacts tracking stability and loop closure success. Conversely, poor descriptors often lead to mismatches, resulting in pose estimation errors and degraded map consistency. Recent advances in self-supervised feature learning, such as SuperPoint [22], have demonstrated the effectiveness of deep neural networks trained on large datasets for extracting robust keypoints in challenging scenarios. However, directly applying learning-based feature extractors, such as SuperPoint, to thermal images leads to cumulative localization errors and frequent tracking failures due to the limited number of reliable keypoints. Our experiments (see Fig. 5) further reveal that SuperPoint produces feature matches in the thermal domain with limited robustness.

To address this, we introduce a dedicated solution for thermal feature learning, termed the Self-supervised Thermal Point (STP) network. As illustrated in Fig. 3, STP leverages descriptor distances between homography-transformed thermal image pairs as self-supervised training signals, enabling effective adaptation to the thermal domain. To address the scarcity of thermal training data, STP is initialized with weights from the SuperPoint, thereby transferring geometric priors learned in the RGB domain to thermal feature learning. **STP Network Architecture** Following SuperPoint architecture, our network consists of a shared encoder for high-level feature abstraction and two decoders, as illustrated in Fig. 3.

The detection head predicts the spatial probability distribution of interest points, producing a tensor  $\mathcal{C} \in \mathbb{R}^{65 \times \frac{H}{8} \times \frac{W}{8}}$ , where the 65 channels represent an  $8 \times 8$  grid cell plus an additional dustbin channel. Reshaping this tensor yields an  $H \times W$  heat map of pixel-level saliency. The descriptor head generates  $\mathcal{D} \in \mathbb{R}^{256 \times \frac{H}{8} \times \frac{W}{8}}$ , which is upsampled via bicubic interpolation to  $\mathbb{R}^{256 \times H \times W}$ . Each descriptor vector is then L2-normalized to unit length. After filtering the detected feature points, their corresponding descriptors are extracted.

**Thermal Homography Training Loss** We freeze the shared encoder and detection head to preserve keypoint localization, while training only the descriptor head to enhance feature representation in thermal datasets. This design obviates the need for pseudo-labels. For each input image, a second view is generated by sampling a homography transform with independent translation, rotation, scaling, and perspective components, ensuring a patch overlap ratio above 0.85. Photometric augmentations are applied to both views, including random brightness and contrast adjustments, additive Gaussian noise, speckle noise, and motion blur.

Given a pair of gray-scale images related by a homography  $H$ , the grid coordinates are  $i = (h_i, w_i) \in \{0, \dots, H_c - 1\} \times \{0, \dots, W_c - 1\}$ ,  $j = (h_j, w_j)$ , where  $H_c = \frac{H}{8}$  and  $W_c = \frac{W}{8}$  in all our experimental setting. We defined  $\mathbf{d}_i \in \mathbb{R}^{256}$  and  $\mathbf{d}'_j$  denote the L2-normalized descriptors at the centers of cell  $i$  in the reference image and cell  $j$  in the warped image, respectively.  $\Delta_{ij}$  is the Euclidean distance between the two cell centres after alignment by  $H$ . We use the following constants  $r_+ = 8$  px,  $m_p = 1$ ,  $m_n = 0.2$ ,  $\lambda_d = 250$ . Then we calculate the positive and negative correspondence matrix by  $s_{ij} = 1$  if  $\Delta_{ij} \leq r_+$ , and 0 otherwise. The similarity is represented as  $s_{ij} \in \{0, 1\}$  and  $c_{ij} = \max(0, \mathbf{d}_i^\top \mathbf{d}'_j) \in [0, 1]$ . The descriptor loss function is defined as

$$\begin{aligned} \mathcal{L}_{ij} &= \lambda_d s_{ij} (m_p - c_{ij})_+ + (1 - s_{ij}) (c_{ij} - m_n)_+, \\ \mathcal{L}_d &= \frac{1}{N} \sum_{i=1}^{H_c W_c} \sum_{j=1}^{H_c W_c} \mathcal{L}_{ij}. \end{aligned} \quad (2)$$

with  $(x)_+ = \max(0, x)$ ,  $N = B(H_c W_c)^2$  denoting the total number of candidate pairs where  $B$  denotes batch size.



is initialized from Eq. 4 and inserted into the map with an observation edge linked to the new KF.

Local BA is asynchronously performed over a sliding window consisting of the current KF and its top- $k$  covisible KFs:

$$\{\mathbf{T}, \mathbf{p}\}^* = \arg \min \sum_{(\text{KF}, i)} \rho(\|\mathbf{u}_i^{\text{KF}} - \pi(\mathbf{T}_w^{\text{KF}} \mathbf{p}_i^w)\|^2). \quad (6)$$

where  $\rho(\cdot)$  denotes a robust loss function. Variables associated with outer KFs are Schur-marginalized, and the reduced system is solved using the Levenberg–Marquardt algorithm.

#### F. Incremental BoW Construction for Thermal Loop-Closing

**Incremental Construction of Binary BoW** Traditional BoW models, such as DBoW2 [27], require extensive offline training and lack cross-modal generalization. This limitation is particularly severe in the thermal domain, where cameras with different specifications often produce images with inconsistent feature distributions. To address this issue, we adopt an iBoW strategy [28], in which the visual dictionary grows online as new descriptors are observed.

For each incoming frame  $t$  with descriptors  $\mathbf{D}_t = \mathbf{d}_i^t$ , each descriptor  $\mathbf{d}_i^t$  is assigned to the nearest visual word

$$\mathbf{w}^* = \arg \min_{\mathbf{w}_m \in \mathcal{V}} \text{Ham}(\hat{\mathbf{d}}_i^t, \mathbf{w}_m) \quad (7)$$

if the Hamming distance is below a threshold  $\tau$ . Otherwise, a new word is created from  $\hat{\mathbf{d}}_i^t$ . In this way, the vocabulary  $\mathcal{V}$  is incrementally constructed from the test sequence itself, eliminating the need for offline training.

Since deep descriptors are high-dimensional feature vectors, we use binary hashing to enable efficient Hamming distance search. Each descriptor component  $d_i$  is binarized as

$$\hat{d}_i = \begin{cases} 1 & d_i \geq 0, \\ 0 & d_i < 0, \end{cases} \quad i = 1, \dots, 256, \quad (8)$$

where  $\hat{\mathbf{d}}$  is the resulting binarized descriptor. This simple yet effective transformation preserves the discriminative structure of the feature space while making descriptors compatible with iBoW indexing.

**Thermal Loop-Closure Detection** To mitigate long-term drift and maintain global consistency, we employ a two-stage loop closure pipeline consisting of appearance-based retrieval and geometric verification. Given the binarized descriptors  $\hat{\mathbf{D}}_t = \{\hat{\mathbf{d}}_i^t\}$  from the current frame, the iBoW index retrieves a ranked set of candidate keyframes (KFs). The similarity score between frame  $t$  and candidate  $j$  is defined as

$$s_{tj} = \frac{1}{|\mathcal{M}_{tj}|} \sum_{(i,k) \in \mathcal{M}_{tj}} (1 - \frac{1}{L} \text{Ham}(\hat{\mathbf{d}}_i^t, \hat{\mathbf{d}}_k^j)), \quad (9)$$

where  $\mathcal{M}_{tj}$  is the set of tentative matches and  $L$  is the descriptor length. To improve temporal robustness, candidates are grouped into islands  $\mathcal{I}$  of size  $L$ , and the best hypothesis is selected as

$$\mathcal{I}^* = \arg \max_{\mathcal{I}} \frac{1}{|\mathcal{I}|} \sum_{j \in \mathcal{I}} s_{tj}. \quad (10)$$

The best candidate KF  $k^*$  undergoes geometric validation by estimating the fundamental matrix  $\mathbf{F}$  with RANSAC and counting inliers consistent with epipolar geometry. If the inliers number exceeds  $\tau_{\text{inl}}$ , a 7-DoF similarity transform  $\mathbf{S} = \{s\mathbf{R}|\mathbf{t}\} \in \text{Sim}(3)$  is computed such that  $\mathbf{x}_i^t \simeq s\mathbf{R}\mathbf{x}_i^{k^*} + \mathbf{t}$ , where  $s \in \mathbb{R}^+$ ,  $\mathbf{R} \in SO(3)$ , and  $\mathbf{t} \in \mathbb{R}^3$ . This transform is refined through non-linear optimization by minimizing the reprojection error. After refinement, correspondences are re-evaluated, and the number of consistent inliers is checked. If the inlier count exceeds a threshold, the loop closure is accepted into the pose graph; otherwise, the hypothesis is discarded.

**Global Pose Optimization** Once validated, the loop adds a new constraint to the pose graph  $\mathcal{G} = (\mathcal{K}, \mathcal{E})$ . We optimize over  $\text{Sim}(3)$  poses  $\{\mathbf{S}_k\}$  by solving

$$\min_{\{\mathbf{S}_k\}} \sum_{(i,j) \in \mathcal{E}} \rho(\|\log_{\text{Sim}(3)}(\mathbf{S}_{ij}^{-1} \mathbf{S}_i \mathbf{S}_j^{-1})\|^2), \quad (11)$$

where  $\mathbf{S}_{ij}$  is the measured relative  $\text{Sim}(3)$  constraint and  $\rho(\cdot)$  is a robust loss. Finally, the global BA jointly refines all camera poses  $\{\mathbf{T}_k\}$  and landmarks  $\{\mathbf{X}_m\}$ :

$$\min_{\{\mathbf{T}_k\}, \{\mathbf{X}_m\}} \sum_{k,m} \rho(\|\pi(\mathbf{T}_k \mathbf{X}_m) - \mathbf{u}_{km}\|^2), \quad (12)$$

in this expression  $\pi(\cdot)$  is the perspective projection and  $\mathbf{u}_{km}$  the observed keypoint.

## IV. EXPERIMENTS

We conducted extensive experiments to evaluate the LST-SLAM system. Section IV-A provides training details and experiment setups. Section IV-B analyzes the performance of the STP network and feature tracking. Section IV-C compares the overall localization accuracy and robustness of our SLAM system against representative classic and learning-based methods. Section IV-D presents ablation studies to assess each module.

### A. Experimental Setup

**STP Network Training Details** The STP network was trained on a large-scale in-vehicle thermal dataset captured with a MAGNITY camera at  $640 \times 480$  resolution. The dataset covers 11 urban road segments recorded at different times of day, with additional test frames collected from 5 unseen segments. Training was performed using the Adam optimizer [29] with an initial learning rate of  $1 \times 10^{-5}$ . The network was trained for 300 epochs with a batch size of 2, and the learning rate was decayed to 30% of its value at epochs 10, 20, 50, 100, and 150. For each epoch, 10,000 images were randomly sampled from a multi-million-frame corpus. Performance was evaluated on 500 validation images using recall and precision. Training on a single NVIDIA GeForce RTX 3080 required approximately three days.

**Thermal SLAM Evaluation** Thermal SLAM experiments were conducted on two platforms. Most evaluations ran on a personal computer with an AMD Ryzen 9 5900HX CPU and an NVIDIA GeForce RTX 3080 GPU. Due to memory constraints, kilometer-scale evaluations of DROID-SLAM




Fig. 4: Qualitative results of feature extraction and tracking on the M2DGR dataset. Each row visualizes feature matches over continuous video segments sampled at larger frame intervals. The three sequences were chosen to represent different motion directions.

[30] and MAST3R-SLAM [31] were carried out on a server with an Intel Xeon Platinum 8358P CPU and four NVIDIA A800 GPUs (80 GiB VRAM each).

### B. Thermal Feature Extraction Evaluation

**Datasets & Baseline** For robustness evaluation, we employed the M2DGR dataset [32], which provides  $640 \times 512$  thermal streams recorded with a PLUG 617 camera. Four feature extractors were evaluated on identical sequences: STP (ours), SuperPoint [22], ORB [11], and SIFT [10].

**Results & Analysis** We assessed feature robustness in terms of tracking quality using two metrics: (i) the number of matches between consecutive frames, and (ii) the number of inlier matches. Putative matches were obtained via mutual nearest-neighbor descriptor search, and inliers were determined using RANSAC [33] with essential matrix estimation. For evaluation, we randomly sampled 50 subsequences from 3 scenes. Each subsequence starts from a random frame, followed by 100 consecutive frames, producing 100 image pairs with progressively larger viewpoint changes.

As shown in Fig. 5, STP consistently achieves more matches and inliers than all baselines. This result indicates that the extracted thermal feature points are both denser and more robust. Notably, SuperPoint often underperforms traditional methods when frame gaps increase, whereas STP markedly improves SuperPoint’s robustness in the thermal domain. In terms of overall performance curves, our method yields improvements of 35.8%, 47.5%, and 48.4% over SuperPoint, SIFT, and ORB, respectively, in the number of inlier points. These results confirm that self-supervised training on thermal data significantly strengthens feature extraction for thermal SLAM. Visualized matches from STP are presented in Fig. 4.

### C. Thermal SLAM in Large-Scale Dynamic Environments

**Datasets & Baseline** The MS<sup>2</sup> dataset provides synchronized stereo-thermal images from autonomous-driving scenes across multiple times of day, providing a kilometer-scale

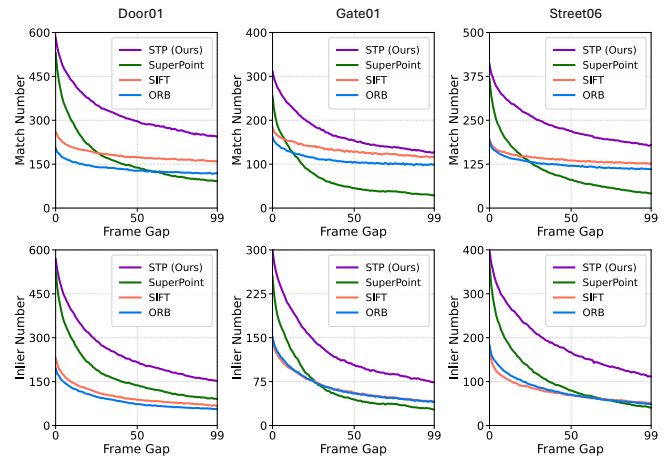


Fig. 5: The STP network significantly outperforms other methods and substantially enhances the SuperPoint behavior in terms of the number of matches and inlier points between different frames.

all-weather benchmark for stereo thermal SLAM [42]. We selected six large-scale sequences with loop closures, collected at different times and under different illumination conditions for the evaluation. Besides, we ran some of the most commonly used and well-known odometry and SLAM systems to compare the localization accuracy and robustness of our LST-SLAM with and without loop closure.

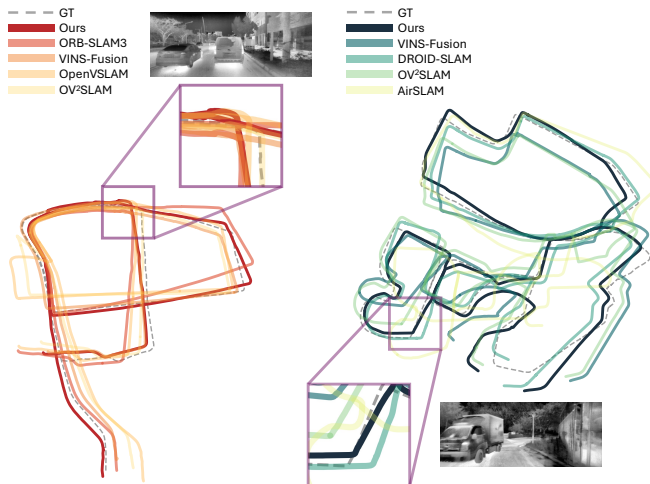
**Results & Analysis** Since the systems incorporate a loop detection component, we adopt the absolute trajectory error (ATE) as the evaluation metric. The root mean square error (RMSE) is computed using  $\text{evo}$  [43]. As this is the first work to evaluate kilometer-scale thermal SLAM datasets, we additionally introduce two metrics:  $t_{\text{apm}} = \text{ATE}_{\text{RMSE}}/L_{\text{gt}}$ ,  $CR = L_{\text{align}}/L_{\text{gt}}$ , where  $L_{\text{gt}}$  is the total ground-truth trajectory length and  $L_{\text{align}}$  is the length of the scale-aligned estimated trajectory overlapping with the ground truth. For each method, we ran 10 iterations and selected the best trajectory for evaluation.

The quantified results are presented in Table I. For  $t_{\text{apm}}$ , we only calculate the average where all sequences are 100%

**TABLE I:** Performance comparison of different odometry and SLAM systems on the MS<sup>2</sup> dataset across various sequences. The average is calculated over all sequences. The best results are in **bold**.

| Sequence                    | Morning      |       |              |       | Daytime      |       |              |       | Nighttime    |       |              |       | Average      |       |
|-----------------------------|--------------|-------|--------------|-------|--------------|-------|--------------|-------|--------------|-------|--------------|-------|--------------|-------|
|                             | Campus       |       | Residential  |       | Campus       |       | Residential  |       | Campus       |       | Residential  |       |              |       |
| Length                      | 7037m        |       | 3090m        |       | 7071m        |       | 3151m        |       | 6970m        |       | 3620m        |       | 5157m        |       |
| Metric                      | $t_{apm}$    | $CR$  | $t_{apm}$    | $CR$  | $t_{apm}$    | $CR$  | $t_{apm}$    | $CR$  | $t_{apm}$    | $CR$  | $t_{apm}$    | $CR$  | $t_{apm}$    | $CR$  |
| <b>Without Loop-Closing</b> |              |       |              |       |              |       |              |       |              |       |              |       |              |       |
| ⊗ SVO 2.0 [34]              | –            | 0%    | –            | 0%    | –            | 0%    | –            | 0%    | –            | 0%    | –            | 0%    | –            | 0%    |
| ⊗ DytanVO [35]              | <b>.0266</b> | 100%  | .0439        | 100%  | .0443        | 100%  | .0559        | 100%  | .0569        | 100%  | .0538        | 100%  | .0469        | 100%  |
| ⊗ TartanVO [36]             | .0379        | 100%  | .0492        | 100%  | .0338        | 100%  | .0497        | 100%  | .0421        | 100%  | .0397        | 100%  | .0421        | 100%  |
| ⊗ <b>LST-SLAM (Ours)</b>    | .0329        | 100%  | <b>.0311</b> | 100%  | <b>.0301</b> | 100%  | <b>.0399</b> | 100%  | <b>.0297</b> | 100%  | <b>.0347</b> | 100%  | <b>.0331</b> | 100%  |
| <b>With Loop-Closing</b>    |              |       |              |       |              |       |              |       |              |       |              |       |              |       |
| ⊗ MAST3R-SLAM [31]          | –            | 5.6%  | –            | 27.7% | –            | 8.9%  | –            | 12.2% | –            | 4.3%  | –            | 17.3% | –            | 12.7% |
| ⊗ ORB-SLAM3 [37]            | –            | 80.3% | –            | 81.3% | –            | 84.1% | .0136        | 100%  | –            | 61.3% | –            | 88.0% | –            | 72.4% |
| ⊗ OpenVSLAM [38]            | –            | 68.3% | .0201        | 100%  | –            | 87.1% | .0179        | 100%  | –            | 28.8% | –            | 84.1% | –            | 77.9% |
| ⊗ OV <sup>2</sup> SLAM [39] | .0157        | 100%  | –            | 61.3% | .0227        | 100%  | .0189        | 100%  | .0276        | 100%  | .0264        | 100%  | .0227        | 93.6% |
| ⊗ AirSLAM [40]              | .0243        | 100%  | .0627        | 100%  | .0222        | 100%  | .0633        | 100%  | .0222        | 100%  | .0583        | 100%  | .0422        | 100%  |
| ⊗ DROID-SLAM [30]           | .0132        | 100%  | .0531        | 100%  | .0146        | 100%  | .0441        | 100%  | .0154        | 100%  | .0438        | 100%  | .0307        | 100%  |
| ⊗ VINS-Fusion [41]          | .0105        | 100%  | .0143        | 100%  | .0166        | 100%  | .0167        | 100%  | .0429        | 100%  | .0132        | 100%  | .0203        | 100%  |
| ⊗ <b>LST-SLAM (Ours)</b>    | <b>.0067</b> | 100%  | <b>.0116</b> | 100%  | <b>.0116</b> | 100%  | <b>.0068</b> | 100%  | <b>.0131</b> | 100%  | <b>.0115</b> | 100%  | <b>.0102</b> | 100%  |

⊗ Monocular input. ⊗ Stereo input.



**Fig. 6:** The qualitative results. **Left:** Daytime Residential. **Right:** Morning Campus. We visualize the trajectories of the top five most accurate methods and the ground truth trajectory for each scene.

complete. For the comparison without loop detection, our method outperforms other VO methods. The SVO results show that traditional monocular odometry methods struggle to initialize in thermal environments. As daytime thermal images resemble grayscale images, our method achieved the highest accuracy in all sequences except for one captured in the morning. The average  $t_{apm}$  of LST-SLAM is 21.4% lower than the second-best system, *i.e.* TartanVO. For the comparison with loop detection, our method exhibits substantially higher robustness than other feature-based SLAM systems such as ORB-SLAM3, OpenVSLAM, and AirSLAM. In particular, compared to AirSLAM, which already demonstrates strong robustness, our approach achieves 75.8% lower localization error. Moreover, relative to the deep learning-based DROID-SLAM and the optical flow-based VINS-Fusion, it achieves 66.8% and 49.8% lower error, respectively. Overall, LST-SLAM outperforms other representative systems under both loop-enabled and loop-disabled conditions. Fig. 6 demonstrates that our method achieves superior pose estimation accuracy and generates smoother trajectories.

**TABLE II:** Performance comparison of different ablation setups. We present the average results of six sequences in MS<sup>2</sup> datasets.

| Method          | Campus       |             | Residential  |             | Average      |             |
|-----------------|--------------|-------------|--------------|-------------|--------------|-------------|
|                 | $t_{apm}$    | $CR$        | $t_{apm}$    | $CR$        | $t_{apm}$    | $CR$        |
| w/o STP         | –            | 76.8%       | .0184        | 100%        | –            | 88.4%       |
| w/o Loop        | .0309        | 100%        | .0352        | 100%        | .0331        | 100%        |
| w/o Seg         | .0140        | 100%        | .0159        | 100%        | .0149        | 100%        |
| w/o DT          | .0127        | 100%        | .0144        | 100%        | .0136        | 100%        |
| <b>LST-SLAM</b> | <b>.0105</b> | <b>100%</b> | <b>.0099</b> | <b>100%</b> | <b>.0102</b> | <b>100%</b> |

#### D. Ablation Study

We evaluated the contribution of each module using four variants:

- *w/o STP*: replaces STP with SuperPoint for thermal feature extraction;
- *w/o Loop*: removes loop closure and global optimization;
- *w/o Seg*: excludes the dynamic object filtering module;
- *w/o DT*: replaces dual-level tracking with the motion model and BoW-guided projection from ORB-SLAM3.

Results, averaged over three time periods for Campus and Residential sequences, are shown in Table II. Directly substituting SuperPoint significantly degrades robustness, preventing the successful evaluation of accuracy. Compared with *w/o Loop*, *w/o Seg*, and *w/o DT*, the full LST-SLAM achieves 69.2%, 31.5%, and 25.0% lower localization error, respectively, demonstrating the effectiveness of each module.

## V. CONCLUSION

This work proposes a stereo thermal SLAM system designed for localization and mapping in dynamic, kilometer-scale thermal environments. It leverages self-supervised learning for thermal feature extraction and employs dual-level tracking for robust feature matching. Together with online iBoW loop detection and global optimization, our approach achieves superior localization accuracy and robustness compared to existing representative VO and SLAM systems. In future work, we aim to further improve the efficiency of LST-SLAM by developing a more lightweight thermal feature extractor. We also plan to explore deep

learning-based feature matching strategies to enhance system robustness and to extend the framework with additional sensing modalities such as IMUs.

## REFERENCES

- [1] A. Macario Barros, M. Michel, Y. Moline, G. Corre, and F. Carrel, "A comprehensive survey of visual slam algorithms," *Robotics*, vol. 11, no. 1, p. 24, 2022.
- [2] S. Vidas and S. Sridharan, "Hand-held monocular slam in thermal-infrared," in *2012 12th International Conference on Control Automation Robotics & Vision (ICARCV)*, pp. 859–864, IEEE, 2012.
- [3] C. Keil, A. Gupta, P. Kaveti, and H. Singh, "Towards long term slam on thermal imagery," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 10174–10181, IEEE, 2024.
- [4] Y. Wu, L. Wang, L. Zhang, Y. Bai, Y. Cai, S. Wang, and Y. Li, "Improving autonomous detection in dynamic environments with robust monocular thermal slam system," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 203, pp. 265–284, 2023.
- [5] P. V. K. Borges and S. Vidas, "Practical infrared visual odometry," *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 8, pp. 2205–2213, 2016.
- [6] W. Chen, Y. Wang, H. Chen, and Y. Liu, "Eil-slam: depth-enhanced edge-based infrared-lidar slam," *Journal of Field Robotics*, vol. 39, no. 2, pp. 117–130, 2022.
- [7] L. Qin, C. Wu, X. Kong, Y. You, and Z. Zhao, "Bvt-slam: a binocular visible-thermal sensors slam system in low-light environments," *IEEE Sensors Journal*, vol. 24, no. 7, pp. 11599–11609, 2023.
- [8] S. Zhao, P. Wang, H. Zhang, Z. Fang, and S. Scherer, "Tp-tio: A robust thermal-inertial odometry with deep thermalpoint," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4505–4512, IEEE, 2020.
- [9] T. Mouats, N. Aouf, D. Nam, and S. Vidas, "Performance evaluation of feature detectors and descriptors beyond the visible," *Journal of Intelligent & Robotic Systems*, vol. 92, no. 1, pp. 33–63, 2018.
- [10] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.
- [11] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "Orb: An efficient alternative to sift or surf," in *2011 International conference on computer vision*, pp. 2564–2571, Ieee, 2011.
- [12] K. M. Yi, E. Trulls, V. Lepetit, and P. Fua, "Lift: Learned invariant feature transform," in *European conference on computer vision*, pp. 467–483, Springer, 2016.
- [13] Y. Lu and G. Lu, "Suprathermal: Matching thermal as visible through thermal feature exploration," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2690–2697, 2021.
- [14] B. Deshpande, S. Hanamsheth, Y. Lu, and G. Lu, "Matching as color images: Thermal image local feature detection and description," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1905–1909, IEEE, 2021.
- [15] Ö. Tuzcuoğlu, A. Köksal, B. Sofu, S. Kalkan, and A. A. Alatan, "Xoftr: Cross-modal feature matching transformer," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 4275–4286, 2024.
- [16] Y.-S. Shin and A. Kim, "Sparse depth enhanced direct thermal-infrared slam beyond the visible spectrum," *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 2918–2925, 2019.
- [17] M. R. U. Saputra, C. X. Lu, P. P. B. de Gusmao, B. Wang, A. Markham, and N. Trigoni, "Graph-based thermal-inertial slam with probabilistic neural networks," *IEEE Transactions on Robotics*, vol. 38, no. 3, pp. 1875–1893, 2021.
- [18] J. Jiang, X. Chen, W. Dai, Z. Gao, and Y. Zhang, "Thermal-inertial slam for the environments with challenging illumination," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 8767–8774, 2022.
- [19] L. Chen, L. Sun, T. Yang, L. Fan, K. Huang, and Z. Xuanyuan, "Rgb-t slam: A flexible slam framework by combining appearance and thermal information," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5682–5687, IEEE, 2017.
- [20] Y. Xu, Q. Hao, L. Zhang, J. Mao, X. He, W. Wu, and C. Chen, "Slam in the dark: Self-supervised learning of pose, depth and loop-closure from thermal images," *arXiv preprint arXiv:2502.18932*, 2025.
- [21] Y. Wu, L. Wang, L. Zhang, M. Chen, W. Zhao, D. Zheng, and Y. Cai, "Monocular thermal slam with neural radiance fields for 3d scene reconstruction," *Neurocomputing*, vol. 617, p. 129041, 2025.
- [22] D. DeTone, T. Malisiewicz, and A. Rabinovich, "Superpoint: Self-supervised interest point detection and description," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pp. 224–236, 2018.
- [23] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 779–788, 2016.
- [24] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *European conference on computer vision*, pp. 740–755, Springer, 2014.
- [25] C. Forster, M. Pizzoli, and D. Scaramuzza, "Svo: Fast semi-direct monocular visual odometry," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 15–22, IEEE, 2014.
- [26] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard, "g 2 o: A general framework for graph optimization," in *2011 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3607–3613, IEEE, 2011.
- [27] D. Gálvez-López and J. D. Tardós, "Bags of binary words for fast place recognition in image sequences," *IEEE Transactions on Robotics*, vol. 28, pp. 1188–1197, October 2012.
- [28] E. Garcia-Fidalgo and A. Ortiz, "ibow-lcd: An appearance-based loop-closure detection approach using incremental bags of binary words," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3051–3057, 2018.
- [29] D. P. Kingma and J. L. Ba, "Adam: A method for stochastic gradient descent," in *ICLR: international conference on learning representations*, pp. 1–15, 2015.
- [30] Z. Teed and J. Deng, "Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras," *Advances in neural information processing systems*, vol. 34, pp. 16558–16569, 2021.
- [31] R. Murai, E. Dexheimer, and A. J. Davison, "Mast3r-slam: Real-time dense slam with 3d reconstruction priors," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 16695–16705, 2025.
- [32] J. Yin, A. Li, T. Li, W. Yu, and D. Zou, "M2dgr: A multi-sensor and multi-scenario slam dataset for ground robots," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2266–2273, 2021.
- [33] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.
- [34] C. Forster, Z. Zhang, M. Gassner, M. Werlberger, and D. Scaramuzza, "SVO: Semidirect visual odometry for monocular and multicamera systems," *IEEE Transactions on Robotics*, 2017.
- [35] S. Shen, Y. Cai, W. Wang, and S. Scherer, "Dytanvo: Joint refinement of visual odometry and motion segmentation in dynamic environments," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4048–4055, IEEE, 2023.
- [36] W. Wang, Y. Hu, and S. Scherer, "Tartanvo: A generalizable learning-based vo," in *Conference on Robot Learning (CoRL)*, 2020.
- [37] C. Campos, R. Elvira, J. J. G. Rodríguez, J. M. Montiel, and J. D. Tardós, "Orb-slam3: An accurate open-source library for visual, visual-inertial, and multimap slam," *IEEE Transactions on Robotics*, vol. 37, no. 6, pp. 1874–1890, 2021.
- [38] S. Sumikura, M. Shibuya, and K. Sakurada, "Openvslam: A versatile visual slam framework," in *Proceedings of the 27th ACM international conference on multimedia*, pp. 2292–2295, 2019.
- [39] M. Ferrera, A. Eudes, J. Moras, M. Sanfourche, and G. Le Besnerais, "Ov<sup>2</sup>slam: A fully online and versatile visual slam for real-time applications," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1399–1406, 2021.
- [40] K. Xu, Y. Hao, S. Yuan, C. Wang, and L. Xie, "Airslam: An efficient and illumination-robust point-line visual slam system," *IEEE Transactions on Robotics*, 2025.
- [41] T. Qin, S. Cao, J. Pan, and S. Shen, "A general optimization-based framework for global pose estimation with multiple sensors," *arXiv preprint arXiv:1901.03642*, 2019.
- [42] U. Shin, J. Park, and I. S. Kweon, "Deep depth estimation from thermal image," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1043–1053, 2023.
- [43] M. Grupp, "evo: Python package for the evaluation of odometry and slam." <https://github.com/MichaelGrupp/evo>, 2017.