

TP-TIO: A Robust Thermal-Inertial Odometry with Deep ThermalPoint

Shibo Zhao¹, Peng Wang², Hengrui Zhang¹, Zheng Fang², Sebastian Scherer¹

Abstract— To achieve robust motion estimation in visually degraded environments, thermal odometry has been an attraction in the robotics community. However, most thermal odometry methods are purely based on classical feature extractors, which is difficult to establish robust correspondences in successive frames due to sudden photometric changes and large thermal noise. To solve this problem, we propose ThermalPoint, a lightweight feature detection network specifically tailored for producing keypoints on thermal images, providing notable anti-noise improvements compared with other state-of-the-art methods. After that, we combine ThermalPoint with a novel radiometric feature tracking method, which directly makes use of full radiometric data and establishes reliable correspondences between sequential frames. Finally, taking advantage of an optimization-based visual-inertial framework, a deep feature-based thermal-inertial odometry (TP-TIO) framework is proposed and evaluated thoroughly in various visually degraded environments. Experiments show that our method outperforms state-of-the-art visual and laser odometry methods in smoke-filled environments and achieves competitive accuracy in normal environments.

I. INTRODUCTION

Robust and accurate state estimation for micro aerial vehicles (MAVs) is of crucial importance as these versatile robots are taking the place of people to fulfill complex and dangerous missions, such as industrial inspection, remote sensing, search and rescue. Due to their flexibility and ability to minimize risks to humans, there is an increased demand for their ability to achieve perception and localization safely and reliably in more challenging environments, such as fire scenes, subterranean settings, as well as GPS-denied environments. To localize in such environments, visible-light camera sensors seem to be a suitable choice for MAVs because of low power consumption, light weight and affordable price. A number of promising visual odometry and visual-inertial odometry methods have been proposed in recent years. Nevertheless, due to the poorly illuminated conditions and airborne obscuring conditions such as dust, fog and smoke in real world, the data of visible-light cameras can significantly degrade, which makes them unreliable for motion estimation.

Compared with visible-light cameras, thermal cameras are robust to illumination variations and obscuring like dust and fog. Thermal cameras, especially the thermal cameras operating in the long wave infrared spectrum (LWIR), have better spectral penetration properties in the presence of

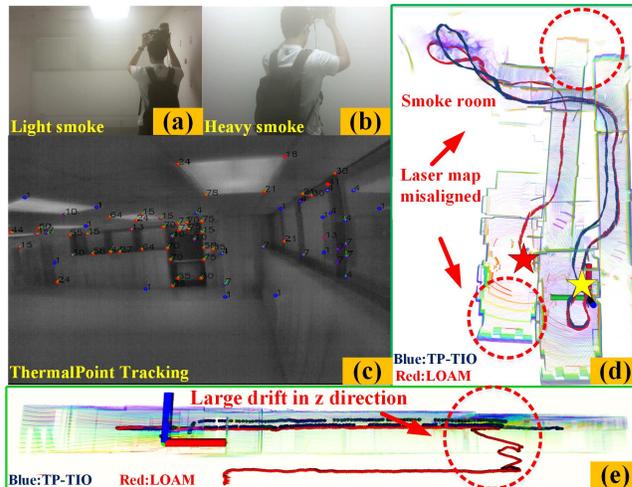


Fig. 1. The perception abilities comparison of different sensors in the same smoke-filled environments. (a) and (b) show the real scene of light-smoke and dense-smoke from a visible-light camera. (c) shows the ThermalPoint tracking results from a thermal camera. Red point means that the feature has been tracked more than ten times. Otherwise, it will be marked as a blue point. (d) shows the trajectory comparison of TP-TIO (blue) and LOAM (red). The red and yellow stars represent the final position of LOAM and TP-TIO, respectively. It can be seen that our method is closer to the origin (xyz axes) while LOAM has a large pose estimation drift and point cloud misalignment. (e) shows the trajectory comparison in the z-direction. It can be seen that our method has a small drift in the z-direction.

fog, dust, and smoke (shown in Fig.1). Therefore, thermal odometry is an ideal choice to achieve robust pose estimation in visually degraded environments. In spite of this, the use cases of the thermal camera are largely limited to perception and localization. Since the thermal camera captures 16-bit temperature radiation images which have low contrast, it is difficult to use handcrafted feature methods to extract salient features. To solve this problem, most of the previous methods [1], [2] re-scale 16-bit thermal images to 8-bit to enhance the contrast of thermal images and improve the performance of feature matching. However, re-scaling thermal image data will result in loss of information, amplifying the effect of noise and image saturation [3]. Moreover, thermal images have fixed-pattern noise and requires a Non-Uniformity Correction (NUC) process [1] to remove noise, which periodically suspends the images for up to 500 milliseconds. This kind of image interruption can easily cause the failure of feature tracking and drift of state estimation.

In the last decade, due to the ability to extract high-level features automatically, deep neural architectures have been proved advantageous for feature extraction in various data types including RGB images and point clouds. However, the shortcomings of deep learning based feature extraction

¹ Shibo Zhao, Hengrui Zhang, Sebastian Scherer are with Robotics Institute, Carnegie Mellon University, USA, {shiboz, hengruiz, basti}@andrew.cmu.edu

² Peng Wang and Zheng Fang are with Faculty of Robot Science and Engineering, Northeastern University, China {pengwang, fangzheng}@mail.neu.edu.cn

networks are also obvious, such as high computation cost and poor real-time performance.

Motivated by the discussion above, we propose a novel deep thermal-inertial odometry (TP-TIO). The main contributions of our work are as follows:

- We propose the first tightly coupled deep thermal-inertial odometry algorithm (TP-TIO), which utilizes both a lightweight learning-based feature detection network and an optimization-based visual-inertial framework.
- We propose the ThermalPoint, a CNN-based feature detection network specifically tailored for producing keypoints on thermal images, that provides notable improvements in anti-noise compared with other state-of-art learning-based and handcrafted feature extraction methods.
- We propose an IMU-aided radiometric feature tracking method, which directly makes use of full radiometric data from an infrared camera to establish robust correspondences between consecutive frames.
- We conducted an extensive set of experiments to thoroughly evaluate our framework from thermal feature tracking to pose estimation in all kinds of environments including obscurant-filled underground mines and dense smoke-filled indoor environments.

The rest of the paper is organized as follows: Section II reviews various feature extraction methods and thermal-inertial odometry methods. Section III illustrates the details of the proposed method. Quantitative and qualitative evaluations from ThermalPoint to state estimation are presented in Section IV. Section V concludes the paper.

II. RELATED WORK

In this section, we start by reviewing existing feature extractor methods, which are essential for vision-based and thermal-based motion estimation, then we review various thermal odometry or thermal SLAM methods.

A. Feature Extractor on Thermal Images

In decades, various handcrafted feature point detection and descriptor algorithms have been proposed and successfully utilized in many visual odometry and SLAM framework. However, since the classical feature extractors are based on gradient information of the images, they cannot work well when images have severe transforms (e.g. viewpoint, blurring, noise). To solve this problem, recent works [4]–[6] leverage deep learning methods to generate more robust interest points and descriptors. Among these methods, LIFT [5] performs both keypoint detection and descriptor extraction based on patch methods and requires a supervision strategy. SuperPoint [4] proposes a homographic adaptation strategy and predicts the keypoints and descriptors from an unsupervised approach. GCNv2 [6] adopts a binary descriptor with good computational efficiency. However, all of these methods are designed for standard visible band images and they are difficult to achieve good performance on thermal imagery.

B. Thermal Odometry

Robust and accurate state estimation from a thermal camera remains a challenging problem. In recent years, there are a number of efforts that have been made for the thermal odometry. To improve the robustness of thermal odometry, many works fuse thermal images with other modalities like RGB images [2], [7], [8] or laser scanner [9]. Since the above methods use traditional handcrafted feature detection methods that mainly rely on the gradient of images, it is difficult for these methods to overcome the effect of large noise. To solve this problem, we propose the "ThermalPoint", which adds a noise reduction module during the training process and effectively overcomes the noise of thermal images. Khatkhat et al. [3] developed a keyframe-based thermal inertial odometry and proved the effectiveness of the direct method. However, since the direct methods heavily rely on image quality and accurate initial pose from the back-end module, it is particularly challenging for these methods to provide robust motion estimation when the environments have a scarcity of thermal gradients. In contrast, since our feature tracking method doesn't rely on the feedback of the back-end module, it can handle temporary feature tracking failure in environments with no salient heat sources. Muhamad [10] proposed an end-to-end thermal inertial odometry. However, this method can only perform well with a slow frame rate. To achieve motion estimation at a high frequency, our method uses ThermalPoint in a shallow architecture and radiometric feature tracking method to establish reliable correspondences instead of matching descriptors. This greatly improves real-time performance and reduces computation.

In summary, it is relatively difficult to find a real-time and robust deep feature-based thermal-inertial odometry under severely visually degraded conditions in the existing literature. This paper aims to solve this problem, providing a reliable sparse feature-based state estimation method in the thermal image domain.

III. THERMAL-INERTIAL ODOMETRY

In this section, we will introduce the pipeline of robust Thermal-Inertial Odometry with deep "ThermalPoint" (TP-TIO) as shown in Fig.2, which allows robust motion estimation in severely visually degraded environments. It is composed of three sequential modules, namely ThermalPoint network, IMU-aided radiometric feature tracking and thermal-inertial bundle adjustment.

A. ThermalPoint Network

1) *Network Architecture*: To detect features robustly on noisy thermal images, we designed a fully convolutional neural network architecture called ThermalPoint shown in Fig.2. Inspired by the famous SuperPoint [4] network, this network has a shared encoder which abstracts high-level features and reduces the dimensionality from the input image. After that, the architecture is divided into two task decoders, feature detection decoder and descriptor decoder, respectively. To improve real-time performance, the descriptor decoder is only used to improve the repeatability [4] of feature detection

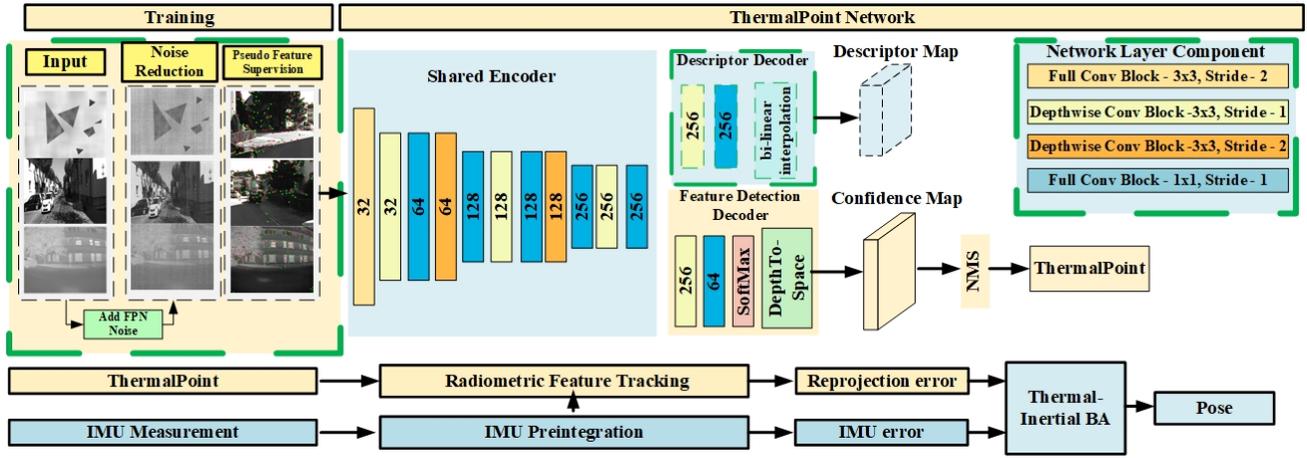


Fig. 2. Overview of deep thermal-inertial odometry. The modules marked as yellow are our contributions to this work. To improve real-time performance, the modules with green dashed squares will only participate in the training process and be excluded in a real application.

during the training process and is removed in the inference process. Instead, we use KLT-based feature tracking method to establish feature correspondences. More details can be found in Section III-C.

2) *Shared Encoder*: Considering the low latency and computation cost for feature detection on embedded systems, our shared encoder is built from an efficient network, MobileNet [11], instead of the VGG-style shared encoder used by SuperPoint. The specific structure of our shared encoder can be found in Fig.2.

3) *Feature Detection Decoder And Descriptor Decoder*: Inspired by SuperPoint [4], the feature detection decoder computes a feature map $X \in \mathbb{R}^{H_c \times W_c \times 64}$. After that, it will output the final confidence heatmap C sized $H \times W$ by performing a channel-wise softmax and DepthToSpace. Each pixel in X corresponds to a probability, measuring a likelihood to be the feature point in the 8×8 cell. For the descriptor part, ThermalPoint is similar to the SuperPoint. Here are some differences we have improved between ThermalPoint and SuperPoint [4].

- To improve the repeatability [4] of feature points, we replace multiple cross-entropy loss used by SuperPoint with an independent logistic loss (Eq.(1)). The reason is that multiple cross-entropy loss only allows one feature extracted in the cell, while independent logistic loss encodes the probability for every pixel.
- To save computation time during training, our decoder performs lower computational bilinear interpolation Y to get dense descriptor map D , instead of bi-cubic interpolation of the SuperPoint.

4) *Loss Function*: Similar to the SuperPoint’s training process, we optimize both the feature point detector loss and the descriptor loss. For the descriptor loss, we directly employ the SuperPoint’s loss. The feature point detector loss function $L_{dec}(I)$ is different, where we use a binary logistic loss over each pixel of C . Given the feature point label

$L(I) \in \mathbb{R}^{H \times W}$ for each image I , the loss can be defined as:

$$L_{dec}(I) = \frac{1}{HW} \sum_{i=1, j=1}^{H, W} -\log(C(i, j))L(i, j) - \log(1 - C(i, j))(1 - L(i, j)) \quad (1)$$

B. Training Details

1) *Noise Reduction*: Applying various noise models to training datasets is a common technique to improve the model’s anti-noise performance. Compared with the noise presented in visual images, thermal images have more complex noises, especially Fixed Pattern Noise (FPN) which is also called non-uniform noise. To overcome the effects of FPN on feature extraction, we need to add FPN noise to the training datasets (Fig.3). Given the non-uniform image $\tilde{I}(i, j)$ and the noise-free image $I(i, j)$, the relation of them can be simply described as the below equation

$$\tilde{I}(i, j) = I(i, j) + n(i, j) \quad (2)$$

where $n(i, j)$ is the non-uniform noise.



Fig. 3. Samples of applying FPN on datasets. The first row shows original noise-free images from synthetic dataset, KITTI dataset and thermal dataset. The bottom row shows corresponding non-uniform images after adding FPN noise.

To obtain non-uniform noise, we employ two simple and effective methods: The first is facing a thermal camera toward a scene with the same material like a concrete wall. In terms of their heat capacity, the non-uniform noise will be recorded due to the scarcity of thermal gradients. The second

is to simulate column FPN. In this paper, we use both the above two methods to create a non-uniform noise dataset which contains 1000 images. During the training process, we randomly select an image from the non-uniform noise dataset and add it to a noise-free image based on Eq.(2) (See Fig.3). Meanwhile, we also use standard data augmentation tricks including contrast, hue, brightness to improve the networks robustness against photometric changes.

2) *Pseudo Feature Supervision*: SuperPoint [4] shows great benefits of using anchors to guide its training process. Although anchors provide a stabler training process, they also prevent the model from identifying new keypoints when no anchor is in the proximity. To solve this problem, we use SIFT [12] to generate some extra pseudo keypoint labels and supervise the training process. By combining the labels from handcrafted features and learned CNN filters, the repeatability of ThermalPoint is improved, as shown in Table I. The specific detail can be found in Fig.4.

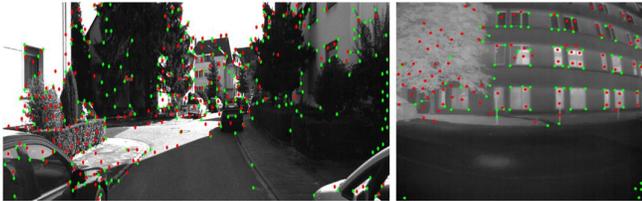


Fig. 4. Pseudo feature supervision training on RGB (left) and Thermal dataset (right), respectively. The green and red points represent keypoints detected by our base feature detector and SIFT [12] method, respectively. Our feature detector mainly detects some "conner" points, while the SIFT method can detect some "salient gradient" points. Both keypoints will be used as labels during the training process.

C. IMU-Aided Radiometric Feature Tracking

After detecting a sparse set of keypoints through our thermal point network, we need to track the motion and build the correspondence of these points between frames. Since most of deep learning-based feature detection networks [4], [5] create dense descriptor map to find correspondences of features, it is difficult to meet our real-time requirements. Inspired by the work [3], instead of using the descriptor method, we use full radiometric based KLT method and IMU measurement to achieve robust feature tracking. To overcome the sudden photometric changes, we combine the inverse-compositional approach [13] with a patch dissimilarity norm, making our feature tracking invariant to intensity scaling.

Given a set of patches in a reference image, we minimize the radiometric error between corresponding patches in two consecutive frames as estimating the transform $\mathbf{T} \in \text{SE}(2)$. For a single residual in a patch, it can be defined as

$$r_i(\xi) = \frac{I_{t+1}(\mathbf{T}\mathbf{x}_i)}{\bar{I}_{t+1}} - \frac{I_t(\mathbf{x}_i)}{\bar{I}_t} \quad \forall \mathbf{x}_i \in \Omega \quad (3)$$

Where $I_t(\mathbf{x})$ is the thermal value of the 16-bit image t at pixel location x . The mean thermal value if the patch in image t is \bar{I}_t . Ω defines the set of image coordinates in the patch.

To improve the robustness for large displacements and real-time performance, we perform our patch alignment over an image scale-space pyramid from coarsest level to finest level, where the alignment results from the upper level provide a prior for the lower level. Furthermore, we use the IMU preintegration [14] to provide transformation prior for the alignment process at the coarsest level. These two steps greatly improve the solution convergence speed and improve the robustness of the tracking.

D. Thermal-Inertial Bundle Adjustment

To estimate the motion of the sensor, an optimization is performed that jointly minimizes visual reprojection error terms of the tracked features and IMU error terms based on IMU preintegration measurements within a sliding window.

1) *Reprojection error*: The first residual term which we use for pose estimation is the reprojection error based on thermal images. Thermal reprojection error can be described as

$$\mathbf{e}_{\text{reproj}} = p^{c_j} - K \left(T_b^c T_w^b l^{w_i} \right) \quad (4)$$

Where p^{c_j} is the point in the target frame j at camera image coordinates c . l^w is the corresponding 3D landmark location in the reference frame i at world coordinates w . T_w^b is the transformation from world frame w to IMU body frame b . T_b^c is the transformation from camera frame c to IMU body frame b .

2) *IMU error*: The second residual term is derived from IMU data based on IMU preintegration measurements. To achieve high frequency and robust pose estimation, we preintegrate several consecutive IMU measurements and add the IMU factor between frame i and frame j . Similar to [14], the corresponding preintegration between IMU body frame b_k and b_{k+1} is $\Delta \mathbf{r} = (\hat{\alpha}_{b_{k+1}}^{b_k}, \hat{\beta}_{b_{k+1}}^{b_k}, \hat{\gamma}_{b_{k+1}}^{b_k})$. Therefore, the error term for inertial measurement within two consecutive frames can be defined as:

$$\mathbf{e}_{\text{imu}} = \begin{bmatrix} \delta \alpha_{b_{k+1}}^{b_k} \\ \delta \beta_{b_{k+1}}^{b_k} \\ \delta \theta_{b_{k+1}}^{b_k} \\ \delta \mathbf{b}_a \\ \delta \mathbf{b}_g \end{bmatrix} \quad (5)$$

$$= \begin{bmatrix} \mathbf{R}_w^{b_k} \left(\mathbf{p}_{b_{k+1}}^w - \mathbf{p}_{b_k}^w + \frac{1}{2} \mathbf{g}^w \Delta t_k^2 - \mathbf{v}_{b_k}^w \Delta t_k \right) - \hat{\alpha}_{b_{k+1}}^{b_k} \\ \mathbf{R}_w^{b_k} \left(\mathbf{v}_{b_{k+1}}^w + \mathbf{g}^w \Delta t_k - \mathbf{v}_{b_k}^w \right) - \hat{\beta}_{b_{k+1}}^{b_k} \\ 2 \left[\mathbf{q}_{b_k}^{w_k^{-1}} \otimes \mathbf{q}_{b_{k+1}}^{w_w} \otimes \left(\hat{\gamma}_{b_{k+1}}^{b_k} \right)^{-1} \right]_{xyz} \\ \mathbf{b}_{ab_{k+1}} - \mathbf{b}_{ab_k} \\ \mathbf{b}_{wb_{k+1}} - \mathbf{b}_{wb_k} \end{bmatrix}$$

Where $(\cdot)^w$ and $(\cdot)^b$ are world frame and IMU body frame respectively. $[\cdot]_{xyz}$ extracts the vector part of quaternions \mathbf{q} . Both rotation matrices \mathbf{R} and Hamilton quaternions \mathbf{q} are used to represent rotation. $\mathbf{p}_{b_k}^w$, $\mathbf{v}_{b_k}^w$ and $\mathbf{q}_{b_k}^w$ are translation, velocity and rotation from the body frame to the world frame while taking the k th image. \mathbf{g}^w is the gravity vector

in the world frame. \mathbf{b}_a and \mathbf{b}_w are accelerometer bias and gyroscope bias respectively. For more specific details about the preintegration theory, we refer the reader to [14].

3) *Joint Optimization*: Since we have obtained the reprojection and inertial residual equations, for each new frame, we minimize the cost function that consists of reprojection residuals, IMU residuals and a marginalization prior \mathbf{e}_m .

$$J = \sum_{i \in \text{obs}(i)} \mathbf{e}_{\text{reproj}}^T W_{\text{reproj}}^{-1} \mathbf{e}_{\text{reproj}} + \sum_{(a,b) \in \mathbf{C}} \mathbf{e}_{\text{imu}}^T W_{\text{imu}}^{-1} \mathbf{e}_{\text{imu}} + \mathbf{e}_m \quad (6)$$

$\text{obs}(i)$ represents a set that contains the ThermalPoint i that is tracked and observed by other frames. The set \mathbf{C} contains pairs of frames (a, b) connected by IMU factors. W_{reproj} and W_{imu} are co-variance matrices for visual and IMU measurements. Through optimizing this cost function, we can obtain the sensor pose.

IV. EXPERIMENTS

To thoroughly evaluate the proposed TP-TIO framework and validate the performance of ThermalPoint, a variety of experiments were performed including underground mine, indoor laboratory and dense smoke-filled environments.

- For the thermal feature detection and tracking, we mainly compare the robustness and real-time performance of ThermalPoint detection and tracking with other state-of-the-art methods including Harris [15], Good feature [16], ORB [17], SIFT [12] and SuperPoint [4].
- For the state estimation, we mainly compare the accuracy and robustness of TP-TIO with other state-of-the-art odometry methods such as VINS [14], OKVIS [18], R-VIO [19] and DSO [20].

The software runs on Jetson AGX Xavier computer with a Linux system running Robot Operating System (ROS). Our experiment video is available at: <https://youtu.be/aa4whgmYTqY>

A. Robustness Comparison of ThermalPoint Detection

To verify the robustness of ThermalPoint detection, we evaluate anti-photometric performance and anti-noise performance. For the anti-photometric test, we use a public underground mine dataset containing severe thermal photometric changes. For the anti-noise test, we use our indoor smoke dataset presenting the challenges of large noise and dense smoke. ThermalPoint is evaluated against other famous state-of-the-art feature detection methods such as Harris, GoodFeature, ORB, SIFT and SuperPoint. We select a suitable fixed threshold for each feature detector to ensure most features can be extracted on this dataset. For this test, we compute a maximum of 400 points for all systems at a 512x640 resolution in each image. The Non-Maximum Suppression(NMS) of all systems is 8.

1) *Anti-photometric and Anti-noise Comparison*: Fig.5 shows the anti-photometric comparison of different feature detectors with sudden photometric changes. We can easily find that deep learning-based methods (ThermalPoint and SuperPoint) can robustly detect enough features even on the low contrast thermal images. However, the classical feature detectors like Harris, Good Feature, ORB and SIFT struggle in the presence of sudden photometric changes. Fig.6 shows the anti-noise comparison of different feature detectors in the presence of fixed-pattern noise appearing as vertical stripes. It can be seen that ThermalPoint extracts few features on the stripe while other methods detect many features on the stripe.

The reason behind this is that our deep-learning-based method can learn anti-photometric and anti-noise abilities through a training process. In contrast, most of the classic detectors are based on the gradient information of images; Therefore, it is difficult for classic detectors to overcome sudden photometric changes.

2) *Quantitative Robustness Analysis*: To further quantitatively analyze the robustness of feature detectors under viewpoint changes and photometric changes on thermal images, we use the repeatability score [4] to evaluate the robustness for each method. If the repeatability score is high, the corresponding method has a good ability of anti-photometric and anti-viewpoint changes. To test the repeatability on thermal images, we apply random homography transformations on our outdoor thermal image sequences to obtain corresponding thermal image pairs. We compute a maximum of 500 points for all systems at a 240x320 resolution. The Non-Maximum Suppression of all systems is 8 and the correctness threshold [4] ϵ is 3. Table I presents the repeatability score for each method with all kinds of changes (viewpoint and photometric). It can be seen that ThermalPoint is the most repeatable under photometric and viewpoint changes.

TABLE I
REPEATABILITY RESULTS OF DIFFERENT FEATURE DETECTORS ON OUR SIMULATED DATASET

Approach	Repeatability
Harris	0.4313
Good feature	0.6253
ORB	0.5748
SIFT	0.4976
SuperPoint	0.7166
Our method	0.7674

B. Robustness and Run-time Comparison of ThermalPoint Tracking

To verify the robustness of ThermalPoint tracking, we test our method in an underground mine environment shown in Fig. 7. During this test, we track a maximum of 160 points for all systems at a 640x512 resolution for every frame. The Non-Maximum Suppression (NMS) of all systems is 8. We find that ThermalPoint can track significantly more points than other algorithms on thermal images with low contrast. Also, we evaluated the real-time performance of various

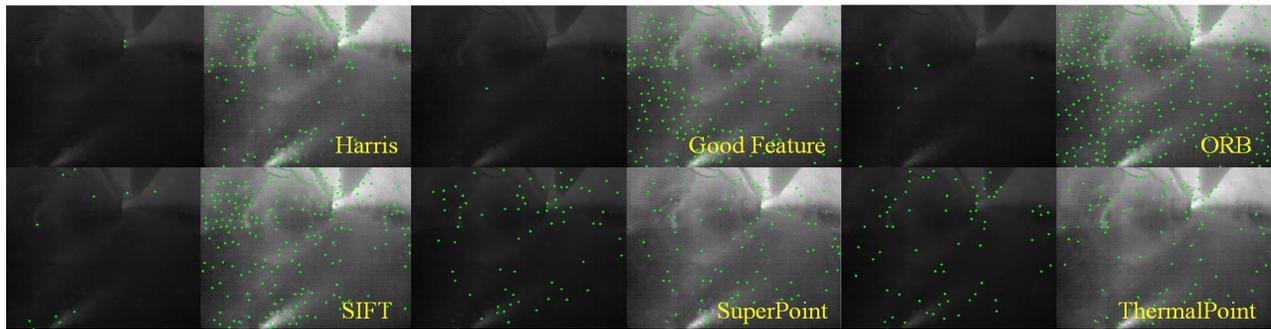


Fig. 5. The anti-photometric comparison of different feature detectors with sudden photometric changes.

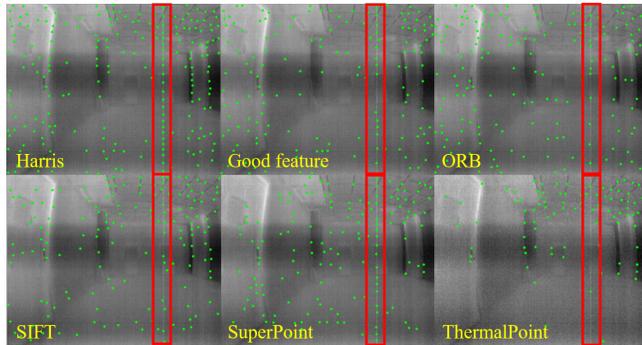


Fig. 6. The anti-noise comparison of different feature detectors in the presence of fixed-pattern noise. The area marked by the red rectangle presents the effect of fixed-pattern noise, which appears as vertical stripes of different brightness in the thermal images.

feature tracking algorithms (Fig. 8). Our feature tracking algorithm takes significantly less time than SuperPoint.

C. Accuracy Comparison of State Estimation in Indoor Environments

To provide a comprehensive accuracy evaluation, the proposed method was compared with four state-of-the-art visual odometry methods, namely VINS [14], OKVIS [25], R-VIO [23] and DSO [19] on thermal imagery. The selections of the above methods are made mainly for two reasons: 1) Most of the thermal odometry approaches [2], [7], [8] require both visual and thermal images to achieve good motion estimation. However, our method works purely on thermal

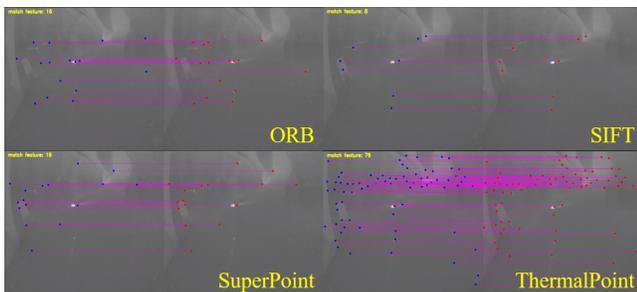


Fig. 7. Samples of different feature tracking method in underground mine dataset.

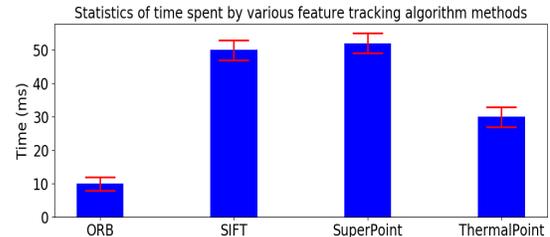


Fig. 8. Average time for detecting and tracking features between two frames

imagery; 2) Since most of the above methods use classic feature detection techniques, by comparing these methods, the effect of feature detection on thermal odometry will be clearly presented.

To achieve a more fair comparison for state-of-the-art methods, we choose our laboratory environment with more heated sources (Fig. 9(a2)) and obtain high contrast thermal images (Fig. 9(a3)), which will be easier for state-of-the-art methods to achieve feature detection and tracking on thermal images. We held our device and walked 85 m along the bridge (Fig. 9(a2)) and returned back to the same position including purely rotational motion. The ground-truth is provided by LOAM [21] algorithm.

Fig. 9(a1,a4) shows the accuracy comparison of estimated trajectories in indoor environments. We can easily find that the trajectory provided by our method (blue line) is the closest to the ground-truth (red line), which suggests that our method is more accurate than other methods. The Absolute Trajectory Error (ATE) and RPE [22] are presented in Table II. It can be noted that the overall ATE and RPE errors of our method are very small.

D. Robustness Comparison of Pose Estimation in Smoke Environments

To evaluate the robustness of our method in challenging environments, we set up smokey environments with a smoke generator. We compare our method with the famous LOAM [21] algorithm to evaluate the accuracy and robustness. The reason for selecting LOAM algorithm is that it can achieve robust pose estimation in light smoke environments while visual odometry will fail in such environments.

In terms of testing the accuracy of trajectories, we use a

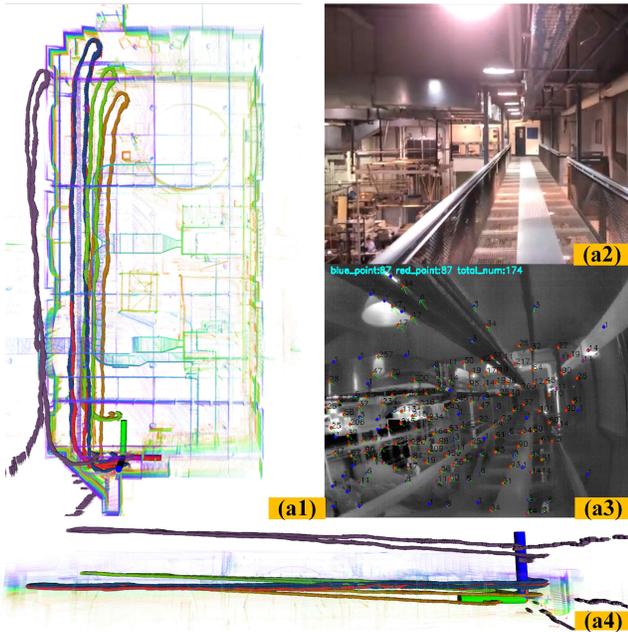


Fig. 9. (a1) and (a4) show the accuracy comparison of estimated trajectories of Ours (blue line), VINS (green line), OKVIS (yellow line) and R-VIO (purple line) against ground-truth (red line) in our laboratory environment as shown in (a2). DSO trajectory was excluded from the plots due to multiple failures. (a3) shows the quality of thermal images and the feature detection result of our method in that scene.

TABLE II
ACCURACY EVALUATIONS OF OURS, VINS, OKVIS, R-VIO AND DSO METHODS OPERATING ON THERMAL IMAGES IN INDOOR ENVIRONMENTS

Approach	ATE (in m)		RPE (in m)	
	RMSE	MAX	RMSE	MAX
OURS	0.359	0.866	0.007	0.153
VINS	2.406	3.902	0.033	0.165
OKVIS	3.199	5.517	0.011	0.258
R-VIO	3.502	4.637	0.029	0.169
DSO	fail	fail	fail	fail

pre-built point cloud map as a ground-truth map to present the estimated trajectory. If the estimated trajectory is accurate, it will coincide with the map well. In this way, we can qualitatively evaluate the accuracy of the trajectory. In terms of the robustness, we use a smoke machine to adjust the amount of smoke and build light smoke environments (Fig. 1(a)) and heavy smoke environments (Fig. 1(b)) to evaluate the performance of the proposed method. In real scenes, we carried our handheld device and walked from room A to the smoke room by the hallway and returned back to the same position in room A as shown in Fig. 10(a).

1) *State estimation comparison under light smoke environments*: In this section, we present the trajectory comparison result between our algorithm (blue line) and LOAM [21] (red line) in light smoke environments. Fig. 10(a1,a2) shows the trajectories crossing a narrow doorway. Our algorithm (blue trajectory) can pass through the doorway while Lidar trajectories fail and collide with the wall, meaning that Lidar

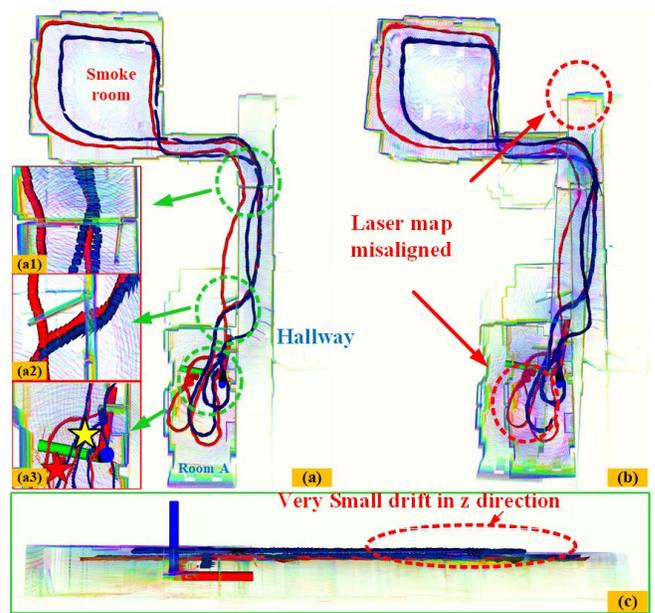


Fig. 10. The state comparison of our method and LOAM under light smoke environments. (a) is ground-truth map and the red and blue trajectories are generated by LOAM and our algorithm respectively. The green circles show more detailed inspections (a1-a3) on the trajectories and map to illustrate the accuracy. (b) shows the actual point cloud misalignment from LOAM. (c) shows estimated trajectories in the pre-built map in z direction.

odometry has a relatively large drift. It's worth noting that the narrow doorway has a width of only 0.67 meters. Fig. 10(a3) shows the result of both trajectories returning to the origin(xyz axes). The red star and yellow star in the figure represent the final position of Lidar odometry and our algorithm, respectively. As we can see, our algorithm's final position is closer to the original start position(xyz axes), indicating that our algorithm is more accurate in this experiment. Fig. 10(b) shows the actual point cloud misalignment from LOAM under light smoke. This suggests that the Lidar point cloud map has a large drift in such environments. Fig. 10(c) shows the trajectories in the pre-built map in the z-direction, and it can be seen that our algorithm has a small z-directional drift.

2) *State Estimation Comparison under Heavy Smoke Environments*: To further test the robustness of our algorithm, we increase the amount of smoke, as shown in Fig. 1(b). Fig. 11(a1,a2) shows the details of trajectories passing through the narrow doorway, and we can see that our odometry (blue line) can pass through the door while the Lidar odometry (red line) fails and collides with the wall. Fig. 11(a3) shows the result of trajectories returning to the origin (xyz axes), where the red and yellow star represent the final position of Lidar odometry and our algorithm respectively. We can see that the final position of our algorithm is closer to the origin, indicating that our trajectory is more accurate in the heavy smoke environment. Fig. 11(b) shows the actual point cloud misalignment from LOAM [21] under heavy smoke. This means that Lidar point cloud map has a large drift in such environments. Fig. 11(c) shows the trajectories in

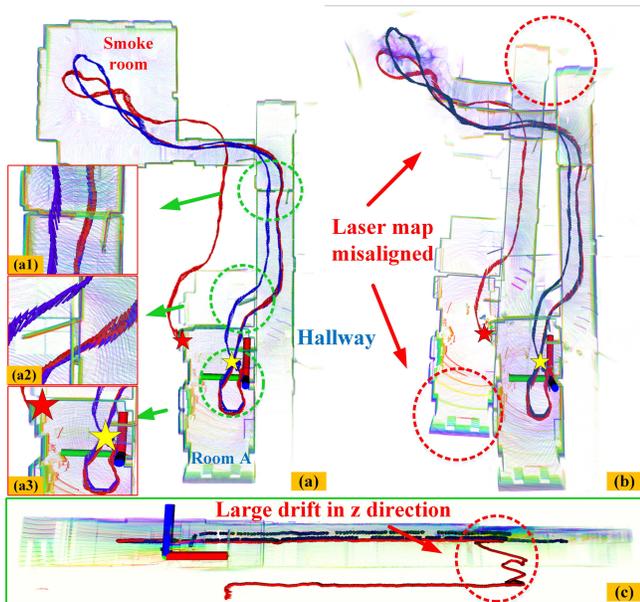


Fig. 11. The state comparison of our method and LOAM under heavy smoke environments. (a) is ground-truth map and the red and blue trajectories are generated by LOAM and our algorithm respectively. The green circles show more detailed inspections (a1-a3) on the trajectories and map to illustrate the accuracy. (b) shows the actual point cloud misalignment from LOAM. (c) shows estimated trajectories in the pre-built map in z-direction.

the pre-built map in z-direction, and it can be seen that our algorithm has small z-directional drift while Lidar odometry (red trajectory) has large drift. This experiment shows that our algorithm is more robust under heavy smoke.

V. CONCLUSION

In this paper, we propose "ThermalPoint", a CNN-based feature detection network specifically tailored for producing keypoints on thermal images. To achieve real-time performance, we combine ThermalPoint with a novel radiometric feature tracking method and establish reliable correspondences between sequential frames. Finally, a deep feature based thermal-inertial odometry (TP-TIO) estimation framework is proposed and evaluated thoroughly in various visually-degraded environments. Experiments show that our method outperforms state-of-the-art visual and laser odometry methods in smoke-filled environments and achieves competitive accuracy in normal environments. In the future, we will add a loop closures module through ThermalPoint and correct accumulated drift for long term navigation.

ACKNOWLEDGMENT

This work was partially supported by DARPA agreement HR00111820044. Also, this work was supported by National Natural Science Foundation of China (No.61573091), the Fundamental Research Funds for the Central Universities (No.N182608003, No.N172608005), Natural Science Foundation of Liaoning Province (20180520006) and Major Special Science and Technology Project of Liaoning Province (No.2019JH1/10100026). We also wish to thank Dr. Shehryar Khattak for his constructive advice.

REFERENCES

- [1] 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.
- [2] T. Mouats, N. Aouf, A. D. Sappa, C. Aguilera, and R. Toledo, "Multispectral stereo odometry," *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, no. 3, pp. 1210–1224, 2014.
- [3] S. Khattak, C. Papachristos, and K. Alexis, "Keyframe-based thermal-inertial odometry," *Journal of Field Robotics*, 2019.
- [4] 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*, 2018, pp. 224–236.
- [5] K. M. Yi, E. Trulls, V. Lepetit, and P. Fua, "Lift: Learned invariant feature transform," in *European Conference on Computer Vision*. Springer, 2016, pp. 467–483.
- [6] J. Tang, L. Ericson, J. Folkesson, and P. Jensfelt, "Gcnv2: Efficient correspondence prediction for real-time slam," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3505–3512, 2019.
- [7] J. Poujol, C. A. Aguilera, E. Danos, B. X. Vintimilla, R. Toledo, and A. D. Sappa, "A visible-thermal fusion based monocular visual odometry," in *Robot 2015: Second Iberian Robotics Conference*. Springer, 2016, pp. 517–528.
- [8] W. Dai, Y. Zhang, D. Sun, N. Hovakimyan, and P. Li, "Multi-spectral visual odometry without explicit stereo matching," in *2019 International Conference on 3D Vision (3DV)*. IEEE, 2019, pp. 443–452.
- [9] 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.
- [10] M. R. U. Saputra, P. P. B. de Gusmao, C. X. Lu, Y. Almalioglu, S. Rosa, C. Chen, J. Wahlstrom, W. Wang, A. Markham, and N. Trigoni, "Deeptio: A deep thermal-inertial odometry with visual hallucination," *IEEE Robotics and Automation Letters*, 2020.
- [11] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "Mobilenets: Efficient convolutional neural networks for mobile vision applications," *arXiv preprint arXiv:1704.04861*, 2017.
- [12] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International journal of computer vision*, vol. 60, no. 2, pp. 91–110, 2004.
- [13] S. Baker and I. Matthews, "Equivalence and efficiency of image alignment algorithms," in *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, vol. 1. IEEE, 2001, pp. I–I.
- [14] T. Qin, P. Li, and S. Shen, "Vins-mono: A robust and versatile monocular visual-inertial state estimator," *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 1004–1020, 2018.
- [15] C. G. Harris, M. Stephens *et al.*, "A combined corner and edge detector," in *Alvey vision conference*, vol. 15, no. 50. Citeseer, 1988, pp. 10–5244.
- [16] J. Shi *et al.*, "Good features to track," in *1994 Proceedings of IEEE conference on computer vision and pattern recognition*. IEEE, 1994, pp. 593–600.
- [17] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "Orb: An efficient alternative to sift or surf," in *2011 International conference on computer vision*. Ieee, 2011, pp. 2564–2571.
- [18] S. Leutenegger, S. Lynen, M. Bosse, R. Siegwart, and P. Furgale, "Keyframe-based visual-inertial odometry using nonlinear optimization," *The International Journal of Robotics Research*, vol. 34, no. 3, pp. 314–334, 2015.
- [19] Z. Huai and G. Huang, "Robocentric visual-inertial odometry," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 6319–6326.
- [20] J. Engel, V. Koltun, and D. Cremers, "Direct sparse odometry," *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 3, pp. 611–625, 2017.
- [21] J. Zhang and S. Singh, "Loam: Lidar odometry and mapping in real-time," in *Robotics: Science and Systems*, vol. 2, no. 9, 2014.
- [22] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of rgb-d slam systems," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 573–580.