

Thermal Homography in Photovoltaic Panels: Evaluating Deep Learning and Feature-Based Methods

Hajer Dahmani*, Mohammed Yaqoob*†, Mohammed Yusuf Ansari†‡, Eduardo Feo Flushing*

*Computer Science, Carnegie Mellon University in Qatar, Doha, Qatar

Email: {hdahmani, ma1, efeoflus}@andrew.cmu.edu

†Electrical and Computer Engineering, Texas A&M University, College Station, TX, USA

Email: ma1@tamu.edu

‡Electrical and Computer Engineering, Texas A&M University at Qatar, Doha, Qatar

Abstract—Thermal imaging has become a vital tool for analyzing temperature variations in various fields, including medical diagnostics, industrial inspection, and environmental monitoring. However, the application of homography techniques in the thermal imaging domain, particularly for photovoltaic (PV) panels, remains underexplored. This paper presents a comprehensive evaluation of state-of-the-art deep homography models, specifically HomographyNet, in comparison with traditional feature-based methods such as ORB+RANSAC and SIFT+RANSAC. These methods are applied to close-range thermal images of PV panels, and their performance is assessed using the mean average corner error (MACE) metric. Extensive experiments analyze the accuracy and robustness of these techniques, as well as the influence of colormap representations on model performance. Key contributions include the evaluation of homography methods for thermal imaging, an in-depth analysis of colormap effects, and the introduction of a novel high-resolution thermal image dataset for PV panels. Results demonstrate that HomographyNet outperforms traditional methods, achieving a MACE of 13.65 with the Noon dataset compared to 24.78 with the Morning dataset. HomographyNets superior performance with magma colormap makes it particularly useful in PV applications such as identifying microcracks or hotspots in solar cells, where accurate thermal image alignment can enhance the visualization of subtle yet critical temperature anomalies.

Index Terms—Thermal imaging, photovoltaic panels, homography, HomographyNet, feature-based methods, colormap effects, deep learning.

I. INTRODUCTION

Imaging technologies have transformed visual data analysis, offering unparalleled precision and detail across various applications [1]. Among these, thermal imaging excels in detecting and visualizing temperature distributions, providing critical insights for medical diagnostics [2], [3], industrial inspection [4], [5], environmental monitoring [6], and surveillance [7]. Although some studies have investigated the integration of thermal imaging with visible light modalities, research on geometric transformations such as perspective correction, homography, and image alignment in the thermal domain is still limited. Photovoltaic (PV) panels, in particular, pose

unique challenges due to their highly reflective surfaces, variable operating conditions, and intricate cell structures. These factors require precise image alignment to accurately detect subtle temperature variations that reveal potential faults or inefficiencies in solar cells [8]. Addressing this gap is crucial for close-range PV panel inspections, as advanced thermal image processing techniques are essential for ensuring the efficiency and reliability of solar energy systems.

Homography, which involves mapping one planar surface onto another, is a fundamental tool for tasks like image stitching [9] and panorama creation [10]. However, applying homography to thermal images presents unique challenges. Unlike visible images that rely on color gradients to define features, thermal images capture infrared radiation patterns linked directly to temperature variations. This reliance on thermal rather than color-based features limits the number of distinct keypoints and often results in reduced texture, low contrast, and fewer identifiable feature points [11]. These constraints are particularly acute in thermal images of PV panels, which tend to exhibit thermal uniformity, minimal distinguishing features, and limited contrast variations [12], [13], all of which complicate reliable feature matching and homography estimation.

This work addresses the need for reliable homography methods tailored to thermal imaging by evaluating both deep learning (DL)-based [14]–[16] and traditional feature-based approaches on thermal images of PV panels. Specifically, we compare HomographyNet [17], a deep homography model, with feature-based methods such as ORB (Oriented FAST and Rotated BRIEF)+RANSAC (Random sample consensus) and SIFT (Scale-Invariant Feature Transform)+RANSAC [18], [19]. We also propose a custom dataset of high-resolution, close-range thermal images of PV panels and analyze each technique using the mean average corner error (MACE) metric [17] across varied inspection scenarios.

Recognizing the limitations in feature richness inherent to thermal images, we also investigate the role of colormap

* Authors Hajer Dahmani and Mohammed Yaqoob have equal contribution.

transformations. Translating temperature data into colormaps enhances homography estimation by improving visual contrast between distinct thermal regions, which in turn boosts keypoint detection and feature matching [20]. To ensure comprehensive evaluation, we tested all colormaps available in the standard Matplotlib library, leveraging their diverse contrast and color schemes to reveal subtle temperature patterns. These colormaps contribute to more robust and reliable homography estimation for thermal images, where contrast and texture are often low.

Our contributions are twofold: First, we provide a comprehensive performance evaluation of DL-based and traditional homography techniques for thermal data, detailing the adaptations necessary to meet the challenges of PV panel inspection. Second, we introduce a novel dataset of thermal images of PV panels, representing diverse capture angles and times of day, designed to reflect real-world inspection conditions. Together, these contributions address a significant gap in thermal image processing, advancing methodologies for precise alignment and analysis in solar energy applications.

II. HOMOGRAPHY ESTIMATION TECHNIQUES

This section reviews homography estimation methods, covering feature-based techniques and advanced DL approaches. It also addresses their core mechanics and challenges in thermal imaging for PV panel inspections.

A. Feature-based Methods

Feature-based methods for homography estimation rely on detecting, describing, and matching keypoints to compute a homography matrix for image alignment. Techniques like ORB and SIFT excel in handling rotation, scale, and illumination variations, while RANSAC refines the homography matrix by minimizing the impact of outliers, ensuring robust alignment accuracy.

The ORB method combines the FAST (Features from Accelerated Segment Test) keypoint detector with the BRIEF (Binary Robust Independent Elementary Features) descriptor, modified for rotation invariance. Keypoints are identified by detecting areas of high contrast, and descriptors are represented as binary strings, making them efficient to compute and match. The matching process typically minimizes a Hamming distance:

$$d_H(\mathbf{b}_1, \mathbf{b}_2) = \sum_{i=1}^n \mathbf{b}_1[i] \oplus \mathbf{b}_2[i],$$

where \mathbf{b}_1 and \mathbf{b}_2 are binary descriptors, and \oplus represents the XOR operation.

SIFT, on the other hand, detects keypoints using the Difference of Gaussian (DoG) approach, where a scale-space representation $L(x, y, \sigma)$ is constructed as:

$$L(x, y, \sigma) = G(x, y, \sigma) * I(x, y),$$

with $G(x, y, \sigma)$ being a Gaussian kernel of scale σ and $I(x, y)$ the input image. The DoG is computed as:

$$D(x, y, \sigma) = L(x, y, k\sigma) - L(x, y, \sigma),$$

where k is a constant multiplicative factor. Keypoints are identified as extrema in the scale-space by comparing $D(x, y, \sigma)$ across scales.

Feature-based methods excel in visible spectrum images with rich textures and distinctive features but face challenges with thermal images. The low contrast and limited texture in thermal images reduce detectable keypoints and matching accuracy. Additionally, the uniform temperature profiles in PV panels hinder feature detection due to the scarcity of distinct gradients required for reliable homography estimation. These challenges highlight the importance of adaptations like colormap transformations to enhance feature detection and improve alignment in the thermal domain.

B. Deep Learning Methods

Deep learning approaches have advanced homography estimation by enabling models to learn transformations directly from image data, bypassing traditional keypoint detection and matching. One notable model in this category is HomographyNet [17], a convolutional neural network (CNN) that estimates relative homography between pairs of images. Unlike conventional methods, HomographyNet operates on stacked grayscale image pairs, directly outputting an 8-degree-of-freedom homography matrix without relying on explicit corner detection or iterative estimation. The network is trained end-to-end on a large dataset of labeled images, allowing it to generalize across various transformations.

While CNN-based models like HomographyNet have demonstrated promising performance on visible light images, their application to thermal images is under-explored. Thermal images capture variations in heat distribution rather than color gradients, leading to notable differences in texture, contrast, and feature visibility compared to visible-spectrum images. These differences introduce unique challenges for CNN-based homography models, as thermal data often lacks the rich gradients and structural details that aid in feature extraction within the visible spectrum. Therefore, assessing the adaptability and effectiveness of DL methods for thermal homography requires targeted evaluation, with possible adaptations such as input transformations or colormaps to enhance contrast and feature definition.

III. DATASET AND METHODOLOGY

This section details the dataset preparation and the implementation of the models.

A. Data Acquisition

Thermal images of PV panels were captured using a FLIR Duo Pro R camera (640x512 resolution) under varying conditions at 8:00 AM and noon to represent colder morning and peak midday temperatures. The images were taken from six distances (10 cm, 15 cm, 20 cm, 25 cm, 30 cm, and 35 cm) at multiple angles (0° and 30°) to introduce variations in scale, proximity, and perspective, enabling a comprehensive evaluation of the performance of homography estimation in various scenarios.

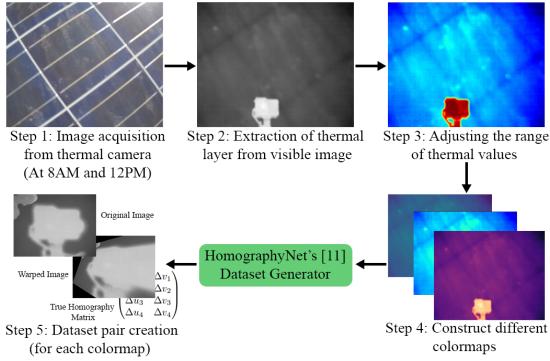


Fig. 1. Overview of the project workflow: from data acquisition to evaluating homography estimation techniques on thermal images.

TABLE I

SIFT AND ORB MODELS AND PARAMETERS (WITHOUT THE DEFAULT MODELS). PATCH SIZE IS REPORTED IN PIXELS; OTHER PARAMETERS DO NOT HAVE UNITS.

Model	Parameter	Value
SIFT Level 1	nfeatures	7000
	contrastThreshold	0.005
	edgeThreshold	5
	sigma	2.6
SIFT Level 2	nfeatures	9000
	contrastThreshold	0.003
	edgeThreshold	7
	sigma	2.8
ORB Level 1	nfeatures	7000
	edgeThreshold	5
	patchSize	20
ORB Level 2	nfeatures	9000
	edgeThreshold	7
	patchSize	40
ORB Level 3	nfeatures	11000
	edgeThreshold	9
	patchSize	60

Applying colormaps is crucial in thermal imaging as they enhance visualization by assigning colors to temperature data, addressing the low contrast and texture inherent in raw thermal images. This improved contrast between thermal regions helps to detect keypoints, facilitating a more accurate homography estimation [20]. We analyzed twenty-two distinct colormaps to evaluate their impact on homography accuracy, identifying those most effective for enhancing thermal data.

Each color-mapped image was converted to grayscale to ensure compatibility across feature-based and DL methods. This conversion preserves the unique intensity gradients and contrast enhancements introduced by each colormap, optimizing feature detection for ORB and SIFT, which rely on intensity gradients for matching accuracy. Additionally, the grayscale images align with HomographyNet's single-channel input requirement, enabling consistent and comparative evaluation across all methods.

Using a dataset generator adapted from [17], we created 710 image pairs with corresponding ground truth homographies, divided into two sub-datasets: Morning (401 pairs) and Noon (309 pairs), as shown in Figure 1. Each sub-dataset captures variations in color mapping, angles, distances, and temperature conditions, enabling a comprehensive evaluation of homography estimation performance in thermal imaging for PV panel inspection.

B. Implementation Details

1) *ORB+RANSAC and SIFT+RANSAC*: The ORB algorithm was implemented using the OpenCV library to detect and match keypoints, with RANSAC employed to estimate the homography matrix by iteratively fitting subsets of matched keypoints and discarding outliers, thereby improving robustness to noise. Initially, ORB was configured with default parameters, followed by fine-tuning across three increasing levels of complexity (Table I). Each level involved adjusting key parameters—specifically, ‘nfeatures’, ‘edgeThreshold’, and ‘patchSize’—to control the number and distribution of detected keypoints. Increasing ‘nfeatures’ at each level allowed for a higher density of keypoints, which improves homography accuracy in images with subtle details. Adjustments to ‘edgeThreshold’ and ‘patchSize’ enabled better control over feature selection in high-contrast and low-contrast regions, respectively, enhancing ORB’s performance across varied colormaps.

Similarly, the SIFT algorithm was fine-tuned across two configurations, with changes made to ‘nfeatures’, ‘contrastThreshold’, ‘edgeThreshold’, and ‘sigma’ parameters. Increasing ‘nfeatures’ and lowering ‘contrastThreshold’ allowed for detecting more fine-grained features, particularly useful in low-contrast thermal images. Adjustments to ‘edgeThreshold’ and ‘sigma’ helped refine the selection of keypoints, improving SIFT’s adaptability to varying thermal textures and enhancing its accuracy in homography estimation. Homography accuracy for both ORB and SIFT was evaluated against ground truth matrices using the MACE metric on 5000 randomly selected image pairs across 22 colormaps, providing a robust comparison of performance under diverse visual conditions.

2) *HomographyNet*: HomographyNet was implemented in TensorFlow and trained in Morning and Noon subdatasets, each divided into 50% training, 25% validation, and 25% testing. Training on an NVIDIA RTX 4090 GPU server ensured sufficient processing power for the large dataset. The model architecture, adapted from [17], features eight convolutional layers with 3x3 filters and ReLU activations, followed by two fully connected layers with 1024 units each. Input images were resized to 128x128 pixels for compatibility and computational efficiency. Training used stochastic gradient descent with a momentum of 0.9 to balance convergence speed and stability, while dropout (0.5) was applied after the final convolutional and first fully connected layers to reduce overfitting. HomographyNet’s performance was assessed using the MACE metric for direct comparison with feature-based methods. The architecture and parameters were adapted from standard practices in homography estimation for visible light images to accommodate the unique characteristics of thermal data.

IV. RESULTS AND DISCUSSION

This section presents the experimental results for both the Morning and Noon datasets, providing a detailed analysis of performance across varying thermal conditions.

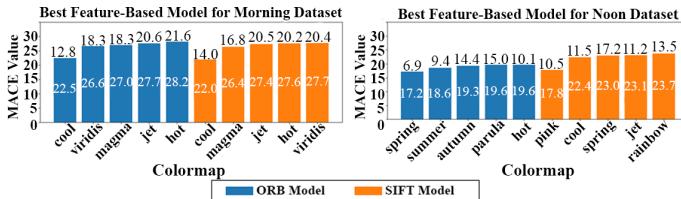


Fig. 2. Mean MACE values achieved by feature-based methods across two datasets, evaluated using various colormaps. The bars represent MACE values, with the corresponding standard deviation values displayed above each bar.

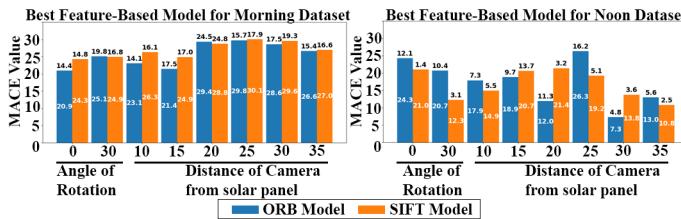


Fig. 3. Performance of Feature-based methods on the two datasets, evaluated for angle of rotation and distance of the image. The bars represent MACE values, with the corresponding standard deviation values displayed above.

A. Feature-Based Model Results

The performance of ORB+RANSAC and SIFT+RANSAC methods was evaluated on the Morning and Noon datasets. Figure 2 presents the top models for five colormaps based on the lowest mean MACE values. In the Morning dataset (left graph), SIFT+RANSAC with the *cool* colormap had a stable mean score of 22.5 (std. 12.8). In contrast, the *hot* colormap showed higher variability with a mean of 28.20 (std. 21.6). *Jet* and *viridis* colormaps had similarly high means and variabilities. ORB+RANSAC followed similar trends, with *cool* yielding the lowest mean and variability, while *jet* and *hot* had the highest means and variabilities, indicating less consistent performance. For the Noon dataset (right graph), the best colormaps for SIFT+RANSAC were *pink*, *cool*, *spring*, *jet*, and *rainbow*. The *pink* colormap had a mean score of 17.83 (std. 10.5), indicating uniform performance. *Jet* and *rainbow* had mean scores around 23, with *rainbow* presenting higher variability (std. 13.52). For ORB+RANSAC, the colormaps *spring*, *summer*, *autumn*, *parula*, and *hot* were evaluated. *Spring* and *summer* had lower mean scores of 17.20 and 18.57 while *Autumn* and *parula* had mean scores around 19, with *parula* showing higher variability (std. 15.05).

Figure 3 shows notable performance variations when evaluating ORB+RANSAC and SIFT+RANSAC techniques on images taken from different angles (0 and 30 degrees) and varying distances from a PV panel. For ORB+RANSAC, the *cool* colormap at 0 degrees yielded a mean score of 20.93 (std. 14.39), whereas the *plasma* colormap at 30 degrees showed a higher mean score of 25.1 (std. 19.81). For different distances, the *cool* colormap at 15 units had a mean score of 21.45, while the *twilight* colormap at 20 units exhibited the highest mean score of 29.37 (std. 24.48). For SIFT+RANSAC, the *cool* colormap at 0 degrees had a mean score of 24.32 (std. 14.84), while the *summer* colormap at 30 degrees resulted in a slightly higher mean score of 24.94 (std. 16.84). Distance-

based evaluation showed the *pink* colormap at 10 units with a mean score of 26.29, and the *autumn* colormap at 25 units with a mean of 30.09.

On the Noon dataset, for ORB+RANSAC, the *turbo* colormap at 0 degrees achieved a mean score of 24.29 (std. 12.13), while the *deepgreen* colormap at 30 degrees had a lower mean score of 20.72 (std. 10.41). At different distances, *turbo* at 30 units had a significantly lower mean score of 7.34 (std. 4.79) while the *hsv* colormap at 25 units scored 26.29 (std. 16.15). For SIFT+RANSAC, the *rainbow* colormap at 0 degrees produced a mean score of 21.04 (std. 1.43), while *viridis* at 30 degrees scored 12.29 (std. 3.09). At different distances, the *turbo* colormap at 20 units scored 21.40 (std. 3.18) and *deepgreen* at 35 units had the lowest mean score of 10.79 (std. 2.53). These results highlight the influence of angles and distances on the performance of both techniques.

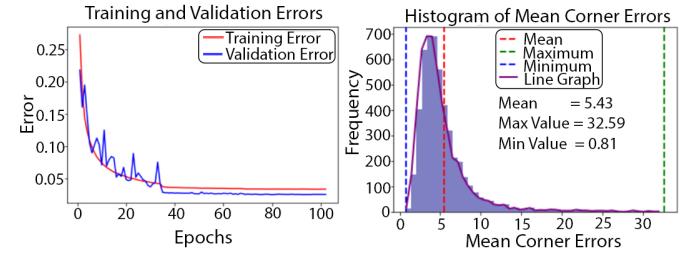


Fig. 4. Performance evaluation of HomographyNet: (left) training and validation loss curves for Magma colormap and (right) histogram results for the same model on the Magma colormap for the Morning dataset.

B. HomographyNet Results

HomographyNet was also evaluated on the two datasets. Figure 4 illustrates the training and validation loss curves (left) and test histogram results (right) for the best-performing model, which utilized the *magma* colormap on the Morning dataset. HomographyNet was tested on the datasets processed with five colormaps: *cool*, *hot*, *jet*, *magma*, and *viridis* (see Figure 5). Each colormap required approximately 8 hours of training, and the *magma* colormap achieved the lowest mean MACE of 5.43. For the Morning dataset, the colormap *cool* achieved a MACE value of 5.89, and *magma* had the lowest MACE value of 5.43. For the Noon dataset, *cool* had a MACE value of 7.11, and *viridis* had the lowest MACE value of 6.39. These results indicate that *viridis* consistently performed best across both datasets, showing the lowest MACE values.

The evaluation of the HomographyNet on the two datasets demonstrates significant variations in MACE values across different angles of rotation and distances from the PV panel (Figure 6). For the Morning dataset, the rotation angle at 0 degrees yielded a MACE value of 3.42, while at 30 degrees, it increased to 6.91. Distance-wise, a camera position at 10 units had a MACE value of 2.13, with the highest value of 6.19 at 30 units. In contrast, the Noon dataset showed more consistent performance, with the rotation angle at 0 degrees, resulting in a MACE value of 2.96 and 3.34 at 30 degrees. The distance from the panel at 10 units had a MACE value of 2.72,

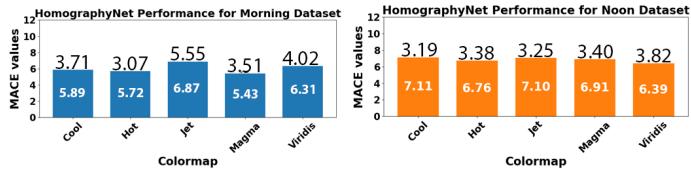


Fig. 5. Mean MACE Values by HomographyNet on different datasets: (left) Morning and (right) Noon, evaluated for various colormaps. The bars represent MACE values, with the corresponding standard deviation values.

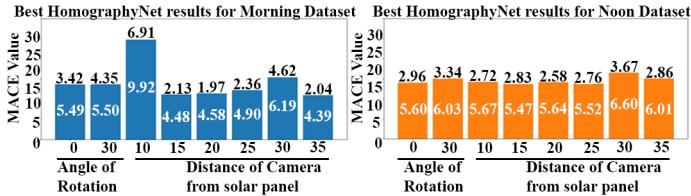


Fig. 6. Performance of HomographyNet on the two datasets, evaluated for angle of rotation and distance of image. The bars represent MACE values, with the corresponding standard deviation values displayed above each bar.

while 35 units showed the highest MACE value of 3.67. The results highlight the variability in HomographyNet's accuracy with changing angles and distances, emphasizing the need for optimal parameter selection for improved performance.

HomographyNet was trained using single-channel raw thermal images, following the same data preparation methodology. Training results indicate that HomographyNet performed better with the Noon dataset, achieving a MACE of 13.65, compared to 24.78 with the Morning dataset. Additionally, the model demonstrated poorer performance on the raw thermal dataset than on the dataset with colormaps. HomographyNet's superior performance with magma colormap makes it particularly useful in PV applications such as identifying microcracks or hotspots in solar cells.

C. Discussion

Our analysis highlights that colormap selection, thermal conditions, and acquisition parameters such as angle and distance heavily influence homography estimation for thermal images. This complexity is critical for PV panel inspection, where subtle thermal variations are key. Both feature-based and DL-based approaches show strengths and limitations, underscoring the need for adaptable strategies to enhance accuracy in thermal imaging.

a) Influence of Colormap Selection on Feature Detection:

A key insight from our results is that colormap selection significantly enhances feature visibility and improves model performance in homography estimation. Thermal images lack the distinct gradients of visible light images, crucial for keypoint detection and matching. By mapping thermal data to specific color palettes, colormaps enhance visual contrast and introduce artificial gradients, enabling feature-based models to detect and match keypoints more effectively. Notably, the *cool* colormap yielded stable, low MACE values across both SIFT and ORB-based models, which we attribute to its capacity to produce balanced contrast without introducing excessive noise. In contrast, colormaps such as *jet* and *hot*, while effective

in certain high-contrast scenarios, often introduced variability due to intense gradients that could either enhance or obscure keypoints depending on the specific thermal patterns present. These results suggest that colormaps with smooth, balanced color transitions are better suited for consistent feature extraction in thermal imaging, especially for scenarios with low inherent contrast like PV panels.

b) Temporal Thermal Variability and Homography Performance:

The observed performance differences between Morning and Noon datasets underscore the need to consider environmental thermal variations when applying homography models to thermal images. During the morning, when temperature gradients across PV panels are relatively low, models often demonstrated higher variability, as evidenced by the performance of *hot* and *viridis* colormaps in SIFT+RANSAC and ORB+RANSAC. The Noon dataset, characterized by more distinct temperature distributions, enabled more stable keypoint detection and matching, particularly with DL-based HomographyNet, which benefited from the well-defined thermal patterns that emerged under peak sunlight conditions. This highlights the model's sensitivity to temporal thermal patterns, suggesting that deploying homography models during periods of maximal contrast can enhance alignment accuracy.

c) Impact of Angle and Distance on Model Robustness:

Variations in angles and distances from the panels further influenced model performance, suggesting that acquisition parameters should be carefully controlled or adjusted based on specific inspection needs. For both ORB+RANSAC and SIFT+RANSAC, greater distances introduced a tendency for higher MACE values and more significant variability, particularly with colormaps like *autumn* and *twilight*, which produced strong gradients but less reliable keypoints at larger scales. This is likely due to the limited resolution of finer features in thermal imaging at a distance, which reduces the precision of keypoint matching. Conversely, close-range images provided better feature clarity, as indicated by the stable MACE values for colormaps such as *cool* and *spring* at shorter distances. HomographyNet also demonstrated sensitivity to these variations, with lower MACE values at 0 degrees and distances under 20 units, highlighting the importance of acquiring images with optimal angles and proximities to maximize model reliability.

d) Deep Learning versus Feature-Based Methods: Adapting to Thermal Imaging Conditions:

While traditional feature-based methods like SIFT and ORB provide adaptability through parameter tuning, our results suggest that they may be less effective under the low-contrast, high-variability conditions typical of thermal imaging without substantial preprocessing through colormaps. DL-based HomographyNet, however, showed a consistent advantage in adapting to colormap-enhanced thermal images, with the *magma* and *viridis* colormaps yielding the lowest MACE values across both datasets. This suggests that DL models may inherently benefit from the spatial patterns introduced by colormaps, learning to identify thermal anomalies and align them with greater robustness. However, HomographyNet's performance declined on raw thermal images without color mapping, underscoring

its reliance on thermal feature richness and the importance of controlled preprocessing for effective generalization.

e) Generalization and Practical Implications: The variability introduced by different colormaps and acquisition settings raises critical questions about the generalizability of homography estimation models in real-world thermal imaging applications. Our findings suggest that no single colormap or acquisition parameter set is universally optimal; instead, the most effective setup is highly dependent on the specific imaging conditions and objectives. For instance, in applications where thermal features are subtle, colormaps like *cool* and *spring* are advantageous, as they provide balanced contrast without obscuring finer details. In contrast, for high-contrast conditions, colormaps such as *hot* or *jet* may offer improved performance but risk greater variability. Thus, selecting and adapting colormaps based on environmental conditions, temporal variability, and acquisition settings is crucial for achieving robust homography estimation in thermal imaging.

V. CONCLUSION

This study underscores the need for adaptive and flexible approaches to homography estimation in thermal imaging, particularly for PV panel inspection, where subtle thermal variations are critical for accurate analysis. Our evaluation reveals that homography performance is closely tied to colormap selection, time of day, and acquisition parameters, all of which must be carefully optimized to achieve reliable results. Traditional feature-based methods like SIFT demonstrate robustness across conditions while fine-tuning ORB models offers improved accuracy under specific scenarios. Colormaps such as *cool*, *spring*, and *summer* were shown to produce consistently stable results, whereas *jet* and *hot* introduce variability, emphasizing the need to tailor colormap selection to enhance feature visibility in thermal images. This adaptability is crucial for both feature-based and DL methods, as optimized colormaps can significantly improve keypoint detection and homography estimation. Notably, HomographyNet, while effective with colormap-enhanced images, exhibited limitations with raw thermal data, suggesting further refinement is necessary for DL models to achieve robustness in diverse thermal imaging conditions. Future work will focus on refining pixel ranges to reduce shadows and sunlight reflections, exploring additional colormaps to optimize homography estimation, and implementing dynamic colormap selection based on real-time thermal conditions to enhance model robustness. Additionally, testing alternative deep learning architectures on the dataset may improve upon HomographyNet, advancing the precision and reliability of homography estimation for PV panel inspection and related applications. Overall, this study encourages adaptive homography techniques in thermal imaging, emphasizing the impact of optimized colormaps and acquisition parameters.

ACKNOWLEDGMENT & CODE

This publication was made possible by the AICC grant AICC04-0715-210006 from the Qatar National Research

Fund. The contents herein reflect the work and are solely the responsibility of the authors.

The code used for benchmarking is available at: [Github Code](#).

REFERENCES

- [1] N. Müller, “Advances in imaging,” *European Respiratory Journal*, vol. 18, no. 5, pp. 867–871, 2001.
- [2] B. F. Jones, “A reappraisal of the use of infrared thermal image analysis in medicine,” *IEEE transactions on medical imaging*, vol. 17, no. 6, pp. 1019–1027, 1998.
- [3] E. Ring and K. Ammer, “Infrared thermal imaging in medicine,” *Physiological measurement*, vol. 33, no. 3, p. R33, 2012.
- [4] S. Taib, M. S. Jadin, and S. Kabir, “Thermal imaging for enhancing inspection reliability: detection and characterization,” *Infrared Thermography*, no. 3, p. 209, 2012.
- [5] K. D. Barnett and M. M. Trivedi, “Analysis of thermal infrared and visual images for industrial inspection tasks,” in *Applications of Artificial Intelligence VII*, vol. 1095. SPIE, 1989, pp. 482–488.
- [6] A. Seward, S. Ashraf, R. Reeves, and C. Bromley, “Improved environmental monitoring of surface geothermal features through comparisons of thermal infrared, satellite remote sensing and terrestrial calorimetry,” *Geothermics*, vol. 73, pp. 60–73, 2018.
- [7] H. Torresan, B. Turgeon, C. Ibarra-Castanedo, P. Hebert, and X. P. Maldaque, “Advanced surveillance systems: combining video and thermal imagery for pedestrian detection,” in *Thermosense XXVI*, vol. 5405. SPIE, 2004, pp. 506–515.
- [8] R. W. Kotla and S. R. Yarlagadda, “Mathematical modelling of spv array by considering the parasitic effects,” *SN Applied Sciences*, vol. 2, no. 1, p. 50, 2020.
- [9] Z. Wang and Z. Yang, “Review on image-stitching techniques,” *Multimedia Systems*, vol. 26, no. 4, pp. 413–430, 2020.
- [10] A. Geetha Kiran and S. Murali, “Automatic rectification of perspective distortion from a single image using plane homography,” *J. Comput. Sci. Appl.*, vol. 3, no. 5, pp. 47–58, 2013.
- [11] M. Vasterling and U. Meyer, “Challenges and opportunities for uav-borne thermal imaging,” *Thermal infrared remote sensing: Sensors, methods, applications*, pp. 69–92, 2013.
- [12] M. Waqar Akram, G. Li, Y. Jin, X. Chen, C. Zhu, X. Zhao, M. Aleem, and A. Ahmad, “Improved outdoor thermography and processing of infrared images for defect detection in pv modules,” *Solar Energy*, vol. 190, pp. 549–560, 2019.
- [13] F. Hong, J. Song, H. Meng, R. Wang, F. Fang, and G. Zhang, “A novel framework on intelligent detection for module defects of pv plant combining the visible and infrared images,” *Solar Energy*, vol. 236, pp. 406–416, 2022.
- [14] M. Yaqoob, M. Ishaq, M. Y. Ansari, V. R. S. Konagandla, T. A. Tamimi, S. Tavani, A. Corradetti, and T. D. Seers, “Geocrack: A high-resolution dataset for segmentation of fracture edges in geological outcrops,” *Scientific Data*, vol. 11, no. 1, pp. 1–13, 2024.
- [15] I. Afsa, M. Y. Ansari, S. Paul, O. Halabi, E. Alataresh, J. Shah, A. Hamze, O. Aboumarzouk, A. Al-Ansari, and S. P. Dakua, “Development and validation of a class imbalance-resilient cardiac arrest prediction framework incorporating multiscale aggregation, ica and explainability,” *IEEE Transactions on Biomedical Engineering*, 2024.
- [16] M. Yaqoob, M. Ishaq, M. Y. Ansari, Y. Qaiser, R. Hussain, H. S. Rabbani, R. J. Garwood, and T. D. Seers, “Advancing paleontology: a survey on deep learning methodologies in fossil image analysis,” *Artificial Intelligence Review*, vol. 58, no. 3, p. 83, 2025.
- [17] D. DeTone, T. Malisiewicz, and A. Rabinovich, “Deep image homography estimation,” *arXiv preprint arXiv:1606.03798*, 2016.
- [18] K. Wu, “Creating panoramic images using orb feature detection and ransac-based image alignment,” *Advances in Computer and Communication*, vol. 4, no. 4, pp. 220–224, 2023.
- [19] R. Redzuwan, N. A. M. Radzi, N. M. Din, and I. S. Mustafa, “Affine versus projective transformation for sift and ransac image matching methods,” in *2015 IEEE International Conference on Signal and Image Processing Applications (ICSIPA)*. IEEE, 2015, pp. 447–451.
- [20] D. Han, J. Shin, H. Kim, and Y. Choi, “The effect analysis of thermal infrared colorization,” in *2021 18th International Conference on Ubiquitous Robots (UR)*, 2021, pp. 1–6.