An Efficient Geometric Transformation-Based Approach for Multi-UAV Image Stitching

Ayoub El-Alami
Laboratory of Modelling and Simulation of Intelligent Industrial Systems (M2S2I), Ecole Normale Superieure de l' Enseignement Technique (ENSET), Mohammedia, Hassan II University of Casablanca (UH2C), Morocco
ayoubalami6@gmail.com (corresponding author)

Younes Nadir
Laboratory of Modelling and Simulation of Intelligent Industrial Systems (M2S2I), Ecole Nationale Superieure de l' Art et de Design (ENSAD), Hassan II University of Casablanca (UH2C), Morocco
younes.nadir@etu.univh2c.ma

Khalifa Mansouri
Laboratory of Modelling and Simulation of Intelligent Industrial Systems (M2S2I), Ecole Normale Superieure de l' Enseignement Technique (ENSET), Mohammedia, Hassan II University of Casablanca (UH2C), Morocco
khalifa.mansouri@enset-media.ac.ma

Received: 25 April 2025 | Revised: 20 May 2025 and 2 June 2025 | Accepted: 8 June 2025

 DOI: https://doi.org/10.48084/etasr.11719

ABSTRACT

UAV-based imaging systems often require real-time image stitching capabilities for tasks such as mapping, surveillance, and environmental monitoring. Traditional approaches typically rely on homography-based transformations, which, although robust, are computationally demanding. These methods often require reestimating transformation parameters for each video frame due to continuous camera motion, significantly increasing computational load and making real-time performance challenging, especially in resource-constrained drone applications. This study proposes an optimized image stitching method based on explicit rotation, scaling, and translation transformations derived from predefined anchor points. Designed for scenarios that involve relatively stable and flat scenes with known anchor markers, this approach removes the reliance on complex matrix operations. Evaluation on experimental footage demonstrates that the method reduces computational complexity, achieving processing times approximately 36% of those required by conventional homography-based methods. Visual assessments and alignment accuracy metrics indicate that the proposed method maintains the same alignment quality, making it particularly suitable for real-time UAV operations and embedded vision applications where computational efficiency is critical.

Keywords-computer vision; image stitching; object detection; object tracking; real-time processing

I. INTRODUCTION

The rapid advancement of drone technology has significantly expanded the potential of aerial imagery for numerous applications, especially in object detection, object tracking, and video surveillance systems [1-4]. Real-time monitoring of dynamic scenes using drone footage presents unique challenges due to limitations in computational resources and the necessity for instantaneous image processing. Effective image stitching is critical in these applications, enabling comprehensive views and continuous tracking across extended fields of view. Real-time image stitching involving multiple Unmanned Aerial Vehicles (UAVs) typically includes several critical steps. Each UAV captures overlapping images to ensure sufficient coverage, which are then preprocessed to standardize parameters such as resolution, brightness, and contrast. Subsequently, distinctive features are identified and matched across images to establish spatial relationships. Geometric transformations are calculated based on these matched features to accurately align the images. After alignment, images are warped onto a common coordinate system and seamlessly blended, minimizing visible seams and color discrepancies. This integrated workflow enables comprehensive and continuous visual coverage, significantly facilitating applications such as surveillance and mapping.

Key steps involved in the stitching process typically include the use of feature descriptors, such as Scale-Invariant Feature Transform (SIFT), Speeded-Up Robust Features (SURF), or Oriented FAST and Rotated BRIEF (ORB), to reliably identify and match distinctive points across images. In addition to these classical feature-based methods, color-based detectors, cascade classifiers (e.g., Haar cascades) [5, 6], and more recent Convolutional Neural Network (CNN) and deep CNN-based detectors have also been employed to enhance feature extraction, particularly under challenging conditions such as low texture or varying illumination [7, 8]. These descriptors enable robust feature matching, facilitating accurate computation of the transformations required for precise image alignment. Feature-based approaches significantly improve the resilience of stitching algorithms to variations in perspective, scale, and illumination, crucial for effective real-time UAV image stitching. Following feature matching, the geometric transformation estimate is calculated to accurately align the images. Subsequently, image warping applies these transformations to position images onto a unified coordinate system. Lastly, blending techniques are employed to merge overlapping regions seamlessly, reducing visual artifacts such as seams and color discrepancies, ensuring cohesive and visually continuous stitched output.

Recent advances in image stitching have presented innovative methods to improve performance, accuracy, and applicability across different domains. In [9], a feature-based stitching algorithm was tailored for microscopic images, focusing on computational efficiency and global alignment, addressing challenges such as uneven illumination and limited overlap. This method demonstrated improved speed and accuracy over existing tools, making it valuable for applications that require rapid processing of high-resolution microscopic data. In [10], an unsupervised image stitching method was based on the YOLOv8 framework, introducing deep homography estimation and attention mechanisms. This approach was designed to improve stitching efficiency and accuracy without relying on supervised learning, making it adaptable to various image stitching scenarios. In [11], a novel method segmented images into multiple regions using superpixel representation and fast density peak clustering, ensuring that the seamline avoids significant ground objects, thus preserving structural integrity in UAV image stitching. In [12], the focus was on automotive applications, introducing a real-time image stitching process that integrates data from multiple vehicle-mounted cameras to create a seamless 360° view, enhancing driver awareness and safety.

Over the past decade, significant advances have been made in image alignment techniques, a crucial step in image stitching. In [13], pixel-wise alignment of overlapping regions was used to guide the adjustment of non-overlapping regions in an unsupervised manner, preserving consistent structure and content for the entire stitched image under limited overlapping conditions. In [14], a method was proposed to warp and align two or more images with overlapping areas to generate a panorama with a larger field of view. In [15], a novel deep image stitching framework was proposed, which exploited the pixel-wise warp field to handle the large-parallax problem, moving beyond traditional homography-based warping. In [16], images were stitched by allowing the alignment of a set of matched dominant planar regions, utilizing rich semantic information directly from the RGB images to extract planar image regions with a deep CNN. In [17], a novel image stitching method used multihomography warping guided by image segmentation to address the large parallax between images. To achieve geometric transformation, traditional methods typically employ homography-based transformations [18-21], which calculate complex matrix operations to align images. Although these methods provide robust alignment, they are computationally intensive, making them less suitable for real-time or embedded applications, such as UAVs deployed in surveillance and tracking tasks. To address these constraints, this study introduces a novel and computationally efficient image stitching approach based on explicit geometric transformations, i.e., rotation, scaling, and translation. The proposed approach leverages predefined anchor points, usually detected by the previous step of feature descriptors, allowing for simplified and direct computation without the need for intensive matrix processing.

The proposed method addresses the prohibitive cost of per-frame homography estimation in real-time UAV stitching by introducing an anchor-based geometric approach. Leveraging two fixed markers, explicit rotation, scaling, and translation parameters are calculated through direct trigonometric relations rather than solving homography matrices for each frame. This simplification reduces the processing time per frame to roughly 35% of that required by conventional estimateAffinePartial2D pipelines while preserving subpixel alignment accuracy, as confirmed by quantitative metrics and visual inspections. Finally, this algorithm was integrated into a lightweight UAV-to-server architecture that offloads most of the computation to a ground station, minimizing resource usage onboard and network latency. Together, these contributions enable efficient high-fidelity stitching in near-planar, marker-instrumented scenes, paving the way for real-time embedded aerial mapping and monitoring [22]. The main contributions of this study are:

- An anchor-based geometric stitching algorithm that replaces matrix-heavy homography solutions with straightforward trigonometric computations.
- A centralized, modular UAV-to-server architecture that offloads stitching to a server node, minimizing onboard processing and network latency.
- A comprehensive evaluation shows about 64% speed improvement without any statistically significant loss in alignment accuracy.

In practice, this method can transform real-time UAV stitching from a computational bottleneck into a low-cost module, accelerating situational awareness in mapping, search-and-rescue, and traffic monitoring missions, where every millisecond of processing time extends battery life and translates directly into safer and more responsive aerial operations. Experimental evaluations illustrate the advantages of the proposed geometric transformation-based method, achieving processing speeds approximately three times faster than conventional homography-based methods while maintaining comparable accuracy. The proposed solution is

particularly effective for scenarios involving relatively flat surfaces, favorable lighting conditions, and clearly visible anchor points, making it exceptionally well-suited for controlled surveillance environments. Future research will focus on adaptive anchor point detection and extend robustness to more challenging conditions, such as varied terrain, complex lighting situations, and partially obscured markers, further expanding the practical utility of the proposed method.

II. THE PROPOSED METHOD

This section presents the proposed image stitching method specifically developed for real-time applications involving UAVs. The primary objective is to significantly reduce computational complexity and runtime while maintaining high-quality image alignment. The proposed method is optimized for scenarios with relatively flat terrain and clearly defined anchor points.

A. Overview

The entire stitching pipeline is orchestrated by a centralized server-side system responsible for coordinating the various modules. Each UAV connects to this server, transmitting its captured image data to extract feature information in real time. The server hosts both the preprocessing and stitching modules, ensuring minimal latency in image alignment and fusion. By centralizing the processing workload, this architecture reduces redundancy across UAVs and allows for more efficient allocation of computational resources, making it especially suitable for embedded and resource-constrained deployments. The proposed drone image stitching pipeline is organized into three main stages, each encapsulated in its own module (Figure 1). First, raw images from each UAV camera are fed into the Preprocessing Module, which contains two independent processors for each drone image stream. After initial operations (edge and anchor detection, perspective rectification), the data passes into the Stitching Module, where geometric parameters are derived, and image warping and fusion take place. The entire process is coordinated by a centralized server. Each UAV connects to this server, transmitting its captured image data and processed anchor information in real time. This centralized architecture ensures that geometric alignment and final stitching can be efficiently handled on the server node with minimal latency. Unlike traditional homography-based stitching techniques, this approach employs a streamlined transformation process that explicitly calculates rotation, scaling, and translation, eliminating the need for computationally intensive calculations and significantly reducing processing time. The method is structured as a modular architecture to enhance flexibility, scalability, and ease of deployment in real-time UAV systems.

B. System Architecture

The system architecture is modular, comprising distinct interconnected components to facilitate a clear division of tasks and efficient data flow. This modularity allows easy scalability, adaptability, and maintenance.

1) UAV Image Acquisition Module

Multiple UAVs simultaneously capture aerial images covering overlapping areas to ensure sufficient spatial overlap. These images are immediately sent to the subsequent processing module for real-time analysis. High-quality and synchronized image capture ensures reliable results throughout the stitching pipeline.

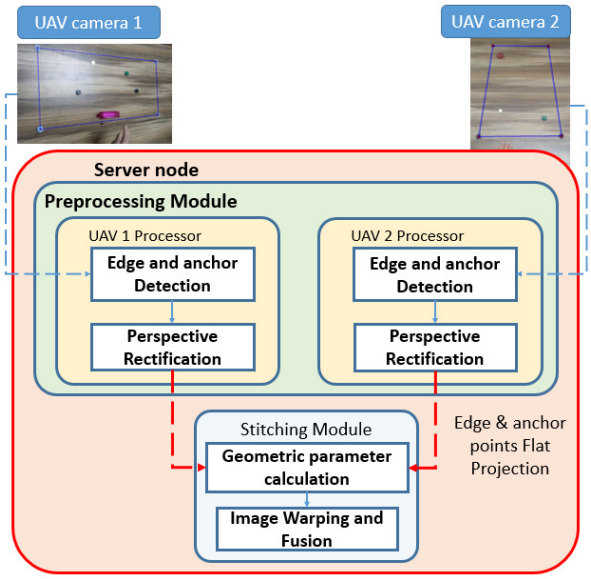

Fig. 1. Architecture of the geometric stitching method.

2) Preprocessing Module

The Preprocessing Module operates in a parallelized manner, with each UAV image stream processed independently by its dedicated processor. This design enables real-time performance and scalability, as multiple UAV inputs can be handled simultaneously without bottlenecking the system. This approach deliberately relies on simple feature descriptors, such as color segmentation, instead of more robust descriptors, such as SIFT or SURF. This design choice aligns with the primary objective of evaluating the performance of the alignment and stitching pipeline itself, independently of the robustness of the feature extraction phase. In a controlled environment, color segmentation is sufficient to accurately detect edges and anchor points. However, in more complex real-world scenarios characterized by challenging textures, lighting variations, or occlusions, robust descriptors would likely be necessary to ensure alignment precision. In particular, the proposed modular system allows seamless integration of such improvements, as the preprocessing module can be easily redefined or extended to incorporate advanced feature extraction techniques when required. Each UAV image undergoes preprocessing involving conversion to the HSV (Hue, Saturation, Value) color space. HSV thresholding isolates specific markers (white and black dots) and edges required for reliable alignment. Morphological operations, including opening and closing, are applied to enhance marker visibility, remove noise, and standardize image quality. This choice is justified by focusing on evaluating the effectiveness of the alignment and stitching process rather than the robustness of the features. Color features are sufficient in this controlled experimental setup to accurately extract edges and anchor points. However, in real-world scenarios that

involve complex textures, lighting variations, or partial occlusions, the use of robust descriptors would be essential to achieve high alignment accuracy. The proposed modular system allows for such future improvements, as the Preprocessing Module can be easily redefined or enhanced to incorporate more advanced feature extraction techniques if needed.

- Anchor Points and Edge Detection: This task focuses on detecting predefined anchor points and edges using color-based segmentation methods. Anchor points are precisely localized using bounding rectangle analysis and subsequently tracked across consecutive frames. Edge detection helps to accurately estimate quadrilateral boundaries, critical for the subsequent rectification process.
- Perspective Rectification: This task performs rectification to project each captured UAV image onto a flat reference plane utilizing the previously identified edges. To achieve this, a simple homography-based transformation maps the active surveilled area, bounded by the four detected edges, onto a predefined flat surface. This transformation ensures that the spatial geometry of the scene is normalized before proceeding to geometric parameter estimation and stitching.

C. Stitching Module

This module receives the rectified images and the final sets of anchor points from both UAV processors. The main steps of this process are shown in Figure 2, and its key roles include determining how to align the two images, computing a global transformation, and blending/merging them into a single image.

1) Geometric Parameter Calculation

This module calculates the transformation parameters, rotation (θ), scale (s), and translation (t_x, t_y), needed to align one image to a designated reference image. Relying on just two matched anchor points avoids the heavier matrix operations of full homography estimation, leading to faster computation. Suppose (x_1, y_1) and (x_2, y_2) are the coordinates of two distinct anchor points (e.g., red and green markers) in the reference image, and (x_1', y_1') and (x_2', y_2') are their corresponding coordinates in the image to be aligned.

Rotation angle θ is given by:

$$\theta = atan2(y_2 - y_1, x_2 - x_1) - atan2(y_2' - y_1', x_2' - x_1')$$

Scale s is given by:

$$s = \frac{\sqrt{(x2 - x1)^2 + (y2 - y1)^2}}{\sqrt{(x2' - x1')^2 + (y2' - y1')^2}}$$

Translation (t_x, t_y) is given by:

$$t_x = x_1 - s * (x_1' * cos(\theta) - y_1' * sin(\theta))$$

$$t_y = y_1 - s * (x_1' * sin(\theta) + y_1' * cos(\theta))$$

These parameters are combined into a single geometric transformation that warps the second image so its anchor points coincide exactly with those of the reference. As illustrated in Figure 3, the colored markers (red and green) are used as anchor points, and the blue boundary denotes the region of interest that defines the transformation bounds. Once the transformation is applied, the second image is aligned on the reference, ensuring a consistent geometric registration.

It is important to note that all steps, from anchor detection through final transformation, are processed on the server side. UAVs only transmit the raw or minimally preprocessed image data and the detected anchor coordinates, thereby reducing on-board computational load and minimizing network latency.

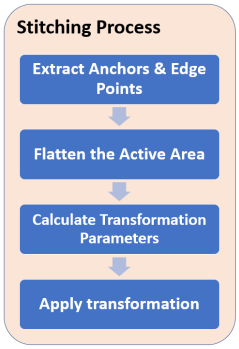

Fig. 2. Main steps of the stitching process.

2) Image Warping and Fusion

Building on the transformation parameters (rotation, scale, and translation) derived in the previous module, this stage applies them to the second image so that its anchor points accurately overlay those of the reference image. First, the warping process utilizes the computed geometric relationships to resample each pixel in the second image, effectively rotating and scaling the image to match the coordinate frame of the reference. This step can be performed through standard image interpolation methods (e.g., bilinear or bicubic), ensuring minimal distortion or artifacts during the alignment process. Once the second image is warped, the system combines (or fuses) the two images into a cohesive mosaic. In the simplest case, fusion may involve directly overlaying the warped image onto the reference, but more advanced approaches can provide a smoother transition. Blending techniques such as alpha masking, feathering, or multiband blending can help reduce visible seams at the overlap boundaries, especially if color or lighting conditions differ between the images. If needed, localized color or brightness correction can be applied before blending to further improve visual consistency. The final result is a seamless stitched image with the anchor points fully aligned, thus accurately capturing a wider scene from two UAV perspectives without noticeable boundaries or distortions. As shown in the experimental image in Figure 4, this step integrates the rectified views of both cameras into a single, uniform coordinate space, enabling the visualization and analysis of the scene as if it were captured by a wider field-of-view camera. In this setup, the green and white markers serve as anchor points, guiding the alignment and ensuring geometric consistency between the two views. In the experimental footage, the left-side image highlights the final image warping and fusion result, where a translucent region denotes the overlapping field of view captured by both UAV cameras. In this overlap zone, where the pink bottle appears, the scene is seen from two slightly different perspectives, but converges

into a single, coherent view in the fused image. This accurate superposition exemplifies the system's ability to handle mutual scene coverage and maintain spatial consistency, confirming that the computed transformation parameters successfully align corresponding objects across the separate camera feeds.

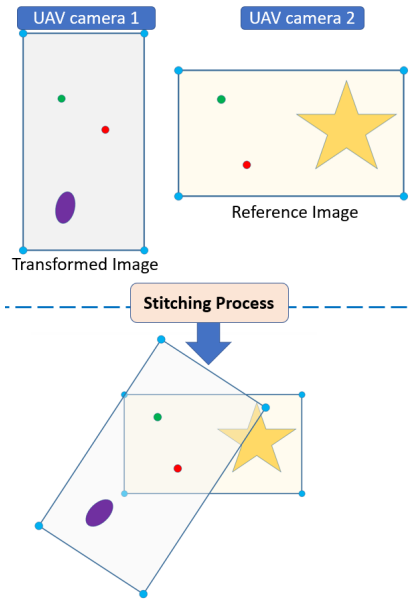

Fig. 3. Illustration of anchor-based alignment and stitching process.

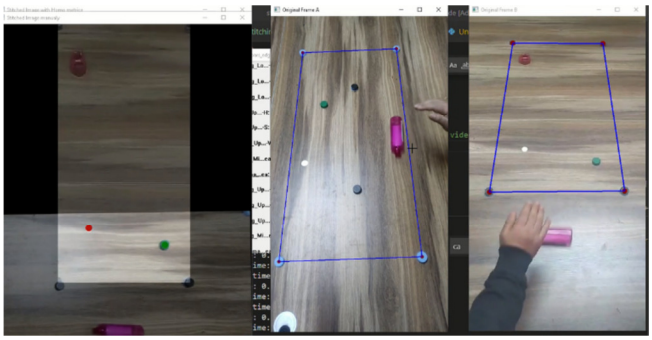

Fig. 4. Experimental setup and output visualization.

III. EXPERIMENT

To validate the effectiveness of the proposed method, experimental evaluations were performed in a controlled laboratory environment designed to mimic realistic UAV surveillance scenarios. The setup involved two cameras capturing different regions of interest from above, with a designated overlapping area that contained clearly visible anchor points and edge markers placed on a flat surface, as shown in Figure 3. Simplified color-based feature descriptors were intentionally used for marker detection to isolate and precisely evaluate the performance of the proposed geometric alignment method. Since the objective of this evaluation was to assess the efficiency of the proposed approach, its processing speed was benchmarked against that of a conventional homography-based stitching method.

A. Experimental Setup

The experiments used two synchronized video streams captured from mounted cameras, each recording at 30 fps with a resolution of 1280×720 pixels. Edge markers (blue dots) were placed to define the boundaries of the active regions of interest, while anchor markers (colored dots) were positioned within these areas to facilitate precise measurement of alignment accuracy. All processing was performed in real time on a machine equipped with an Intel Xeon CPU and 32 GB of RAM.

B. Results and Discussion

The performance of the proposed geometric alignment method was compared against OpenCV's widely used *estimateAffinePartial2D* function, which computes a transformation matrix (covering rotation, translation, and uniform scaling) based on matched points. By design, *estimateAffinePartial2D* often involves more general-purpose matrix operations and optional robust estimators (e.g., RANSAC) for handling outliers. In contrast, the proposed method focuses on a direct lightweight approach that uses only two reliable anchor points. Both methods were tested with identical input data and the same anchor points, allowing a direct comparison of alignment accuracy and computational speed. The stitched outputs of this approach and the OpenCV function exhibited visually and geometrically equivalent alignment when overlaid, confirming the correctness and reliability of the proposed transformation logic. This strongly suggests that the simpler rotation-scaling-translation solution adequately captures the geometric relationships under the test conditions. However, significant differences in processing speed emerged.

Table I summarizes three test scenarios that cover six video sequences with resolutions of 722×406 , 1088×612 to 935×525, and 1408×792 px, to determine whether frame size affects the performance of the proposed method. Across these three tests, the average processing time was nearly constant, driving an average ratio of only 36% compared with OpenCV's estimateAffinePartial2D, approximately 2.8 times faster, while maintaining the same geometric accuracy. These findings validate the effectiveness and efficiency of the proposed method. In scenarios where the scene can be reliably described with only a rotation, scale, and translation model (e.g., near-planar surfaces, consistent camera heights), the proposed method provides comparable alignment at a fraction of the computational cost.

TABLE I. PROCESSING-TIME BENCHMARK ACROSS MULTIPLE VIDEO RESOLUTIONS

			Average processing time (s)		**B/A**
		Dimensions (W×H)	**Traditional method (A)**	**Proposed method (B)**	
Test 1	Sequence 1	(1088×612)	0.014912	0.005476	36.72%
	Sequence 2	(722×406)			
Test 2	Sequence 1	(1280×720)	0.016976	0.006151	36.23%
	Sequence 2	(850×478)			
Test 3	Sequence 1	(1408×792)	0.016245	0.006023	37.08%
	Sequence 2	(935×525)			

This speed improvement can be attributed to:

- Reduced Matrix Complexity: The proposed method avoids matrix inversions and robust estimators, using only basic trigonometric and arithmetic operations.
- Fewer Required Keypoints: The proposed approach explicitly relies on two anchor points, bypassing the need for iterative outlier detection or RANSAC loops.
- Tailored Transformation Model: By limiting the transformation to rotation, uniform scale, and translation, the specific geometry of the UAV scenario is fit without overparameterization.

IV. CONCLUSION

This paper presents an efficient geometric transformation-based approach for real-time image stitching in UAV systems. By focusing on a rotation, scale, and translation alignment strategy, rather than a full homography, this method significantly reduces computational overhead, delivering approximately 65% speed improvement, without sacrificing accuracy. Additionally, the modular system architecture enables scalability, flexibility, and straightforward integration into various application scenarios, from object detection to multi-drone coordination. The method's computational simplicity and lightweight design make it particularly well-suited to real-time applications on resource-constrained UAV platforms, where power and processing limits are critical constraints. Future work will focus on enhancing robustness to complex real-world environments through the incorporation of advanced feature descriptors and adaptive anchor point detection, addressing challenges such as dynamic scenes, variable illumination, and nonplanar terrains. To further validate and extend the proposed approach, additional experiments are planned in outdoor drone datasets featuring heterogeneous backgrounds and moving objects, thus broadening the scope of applicability and pushing the boundaries toward robust, reliable, and highly efficient UAV-based video surveillance systems. A video demonstration of the system is available in [23].

DATA AVAILABILITY

All data, including raw videos and configuration files, along with the source code used, are publicly available in [24].

REFERENCES

[1] A. El-Alami, Y. Nadir, and K. Mansouri, "A review of object detection approaches for traffic surveillance systems," *International Journal of Electrical and Computer Engineering (IJECE)*, vol. 14, no. 5, Oct. 2024, Art. no. 5221, https://doi.org/10.11591/ijece.v14i5.pp5221-5233.

[2] W. Luo, J. Xing, A. Milan, X. Zhang, W. Liu, and T. K. Kim, "Multiple object tracking: A literature review," *Artificial Intelligence*, vol. 293, Apr. 2021, Art. no. 103448, https://doi.org/10.1016/j.artint.2020.103448.

[3] M. S. K. Namana and B. U. Kumar, "Enhancing Surveillance Systems leveraging AIoT for Advanced Object Detection in Real-Time Security Applications," *Engineering, Technology & Applied Science Research*, vol. 15, no. 3, pp. 22507–22517, Jun. 2025, https://doi.org/10.48084/etasr.9926.

[4] A. El-alami, Y. Nadir, and K. Mansouri, "A Modified Lightweight DeepSORT Variant for Vehicle Tracking," *International Journal of Advanced Computer Science and Applications*, vol. 15, no. 10, 2024, https://doi.org/10.14569/IJACSA.2024.0151067.

[5] K. Joshi and M. I. Patel, "Recent advances in local feature detector and descriptor: a literature survey," *International Journal of Multimedia Information Retrieval*, vol. 9, no. 4, pp. 231–247, Dec. 2020, https://doi.org/10.1007/s13735-020-00200-3.

[6] S. K. Sharma, K. Jain, and A. K. Shukla, "A Comparative Analysis of Feature Detectors and Descriptors for Image Stitching," *Applied Sciences*, vol. 13, no. 10, May 2023, Art no. 6015, https://doi.org/10.3390/app13106015.

[7] T. Georgiou, Y. Liu, W. Chen, and M. Lew, "A survey of traditional and deep learning-based feature descriptors for high dimensional data in computer vision," *International Journal of Multimedia Information Retrieval*, vol. 9, no. 3, pp. 135–170, Sep. 2020, https://doi.org/10.1007/s13735-019-00183-w.

[8] H. Kim, W. K. Jung, Y. C. Park, J. W. Lee, and S. H. Ahn, "Broken stitch detection method for sewing operation using CNN feature map and image-processing techniques," *Expert Systems with Applications*, vol. 188, Feb. 2022, Art. no. 116014, https://doi.org/10.1016/j.eswa.2021.116014.

[9] F. S. Mohammadi, H. Shabani, and M. Zarei, "Fast and robust feature-based stitching algorithm for microscopic images," *Scientific Reports*, vol. 14, no. 1, Jun. 2024, Art. no. 13304, https://doi.org/10.1038/s41598-024-61970-y.

[10] C. Qin and X. Ran, "Efficient Unsupervised Image Stitching Using Attention Mechanism with Deep Homography Estimation," *Computers, Materials and Continua*, vol. 79, no. 1, pp. 1319–1334, Apr. 2024, https://doi.org/10.32604/cmc.2024.048850.

[11] W. Pan, A. Li, X. Liu, and Z. Deng, "Unmanned aerial vehicle image stitching based on multi-region segmentation," *IET Image Processing*, vol. 18, no. 14, pp. 4607–4622, Dec. 2024, https://doi.org/10.1049/ipr2.13271.

[12] A. Pujol Miró, "Real-time image stitching for automotive 360º vision systems," B.S. Thesis, Universitat Politècnica de Catalunya, 2014.

[13] Q. Jia, X. Feng, Y. Liu, X. Fan, and L. J. Latecki, "Learning Pixel-wise Alignment for Unsupervised Image Stitching," in *Proceedings of the 31st ACM International Conference on Multimedia*, Jul. 2023, pp. 1392–1400, https://doi.org/10.1145/3581783.3612298.

[14] M. Lin, T. Liu, Y. Li, X. Miao, and C. He, "Image stitching by disparity-guided multi-plane alignment," *Signal Processing*, vol. 197, Aug. 2022, Art. no. 108534, https://doi.org/10.1016/j.sigpro.2022.108534.

[15] H. Kweon, H. Kim, Y. Kang, Y. Yoon, W. Jeong, and K. J. Yoon, "Pixel-wise Deep Image Stitching." arXiv, Dec. 12, 2021, https://doi.org/10.48550/arXiv.2112.06171.

[16] A. Li, J. Guo, and Y. Guo, "Image Stitching Based on Semantic Planar Region Consensus," *IEEE Transactions on Image Processing*, vol. 30, pp. 5545–5558, 2021, https://doi.org/10.1109/TIP.2021.3086079.

[17] T. Liao, C. Wang, L. Li, G. Liu, and N. Li, "Parallax-tolerant Image Stitching via Segmentation-guided Multi-homography Warping." arXiv, 2024, https://doi.org/10.48550/ARXIV.2406.19922.

[18] L. Nie, C. Lin, K. Liao, M. Liu, and Y. Zhao, "A view-free image stitching network based on global homography," *Journal of Visual Communication and Image Representation*, vol. 73, Nov. 2020, Art. no. 102950, https://doi.org/10.1016/j.jvcir.2020.102950.

[19] L. Nie, C. Lin, K. Liao, S. Liu, and Y. Zhao, "Unsupervised Deep Image Stitching: Reconstructing Stitched Features to Images," *IEEE Transactions on Image Processing*, vol. 30, pp. 6184–6197, 2021, https://doi.org/10.1109/TIP.2021.3092828.

[20] S. Bang, H. Kim, and H. Kim, "UAV-based automatic generation of high-resolution panorama at a construction site with a focus on preprocessing for image stitching," *Automation in Construction*, vol. 84, pp. 70–80, Dec. 2017, https://doi.org/10.1016/j.autcon.2017.08.031.

[21] W. Lyu, Z. Zhou, L. Chen, and Y. Zhou, "A survey on image and video stitching," *Virtual Reality & Intelligent Hardware*, vol. 1, no. 1, pp. 55–83, Feb. 2019, https://doi.org/10.3724/SP.J.2096-5796.2018.0008.

[22] M. Fu *et al.*, "Image Stitching Techniques Applied to Plane or 3-D Models: A Review," *IEEE Sensors Journal*, vol. 23, no. 8, pp. 8060–8079, Apr. 2023, https://doi.org/10.1109/JSEN.2023.3251661.

[23] A. El-Alami, "An Efficient geometric based Approach for Real-Time multi UAV Image Stitching," Apr. 03, 2025. https://www.youtube.com/watch?v=7Af9zBixMUA.

[24] A. El-Alami, "Geometric-Stitching." May 20, 2025, [Online]. Available: https://github.com/ayoubalami/Geometric-Stitching.