

**A LEARNING-CENTRIC END-TO-END HYBRID SYSTEM  
FOR IN-SITU THERMOGRAPHIC INSPECTION  
IN AUTOMATED FIBER PLACEMENT**

by  
MUHAMMED ZEMZEMOĞLU

Submitted to the Graduate School of Engineering and Natural Sciences  
in partial fulfilment of  
the requirements for the degree of Doctor of Philosophy

Sabancı University  
July 2025

**A LEARNING-CENTRIC END-TO-END HYBRID SYSTEM  
FOR IN-SITU THERMOGRAPHIC INSPECTION  
IN AUTOMATED FIBER PLACEMENT**

Approved by:

Prof. Dr. MUSTAFA ÜNEL .....  
(Thesis Advisor)

Prof. Dr. BAHATTİN KOÇ .....

Assist. Prof. Dr. MELİH TÜRKSEVEN .....

Prof. Dr. HAYDAR LİVATYALI .....

Assoc. Prof. Dr. ALİ FUAT ERGENÇ .....

Date of Approval: July 22, 2025

MUHAMMED ZEMZEMOĞLU 2025 ©

All Rights Reserved

## ABSTRACT

# A LEARNING-CENTRIC END-TO-END HYBRID SYSTEM FOR IN-SITU THERMOGRAPHIC INSPECTION IN AUTOMATED FIBER PLACEMENT

MUHAMMED ZEMZEMOĞLU

Mechatronics Engineering, Ph.D. Thesis, July 2025

Dissertation Supervisor: Prof. Dr. MUSTAFA ÜNEL

**Keywords:** Automated Fibre Placement (AFP), In-Situ Process Monitoring, Thermographic Inspection, Defect Identification, Quality Assessment, Motion Estimation, Lay-up Reconstruction, Computer Vision, Deep Learning

Automated Fiber Placement (AFP) technology continues to lead and transform composite manufacturing, but its progress remains constrained by persistent quality assurance challenges. Emerging material and process defects compromise structural integrity, while inspection practices remain largely manual, reactive, and incapable of real-time feedback—leading to costly downtimes. Existing methods face two core limitations: partial, mostly offline frame-wise analysis and the lack of global, temporally consistent lay-up visualization. To address these challenges, this thesis proposes a dual-framework, real-time, learning-centric inspection system that unifies local defect intelligence with global lay-up traceability—operating machine-independently using only thermal imagery.

The first framework implements a hybrid, frame-wise defect analysis and quality assessment system composed of three synergistic modules—Dynamic Tow Identification, Hierarchical Defect Identification, and Lay-up Quality Evaluation—configured for parallel or conditional execution to ensure runtime efficiency. It begins with setup-independent spatial-temporal analysis that estimates tow boundaries with sub-pixel accuracy (mean error  $< 0.8$  px), enabling tow-level reasoning. In parallel, high-level defect detection uses a Gabor-based SVM classifier exceeding 95% accuracy and recall. Defective frames are forwarded to a custom 12-layer deep con-

volutional neural network for fine-grained classification, achieving 96.4% multi-class accuracy across defect types. Upon detection, a seeded active contour model adapted to thermal textures yields interpretable segmentation masks with 93.2% pixel accuracy—outperforming baseline methods under noise. These outputs are fused with tow geometry to compute the novel Defect Area Percentage (DAP) metric, which quantifies severity at both tow and course levels and forms the core of the operator decision support system (AFP-DSS). Operating at 5 fps, the framework enables fully autonomous, real-time AFP quality inspection.

The second framework introduces ThermoRAFT-AFP, a machine-independent, deep learning-based thermal motion estimation core tailored to AFP. Built upon the RAFT optical flow, it incorporates AFP-specific augmentations and runtime optimizations (e.g., predictive initialization, drift correction, adaptive early exit) for stable, precise thermal flow tracking. It estimates dense inter-frame motion to drive a two-stage reconstruction pipeline that generates course-wise mosaics and assembles high-fidelity, ply-level laminate visualizations. This restores temporal consistency and global alignment across evolving lay-ups, enabling traceable defect analysis and mirroring industrial inspection workflows. ThermoRAFT-AFP achieves a velocity estimation RMSE of 4.83 mm/s and cumulative drift below 0.1%, while maintaining robustness down to  $\text{SNR} = 14.4 \text{ dB}$  and sustaining real-time operation at 25 fps. Robust against setup variations without retraining and near-zero tuning, it produces interpretable, temporally aligned visualizations that support laminate-scale defect propagation analysis.

Together, the two frameworks form an integrated system that fuses frame-wise defect outputs with global reconstructions into a unified thermal quality view. This fusion links semantic analysis with temporal context in an interpretable visualization aligned with real-world inspection workflows. Evaluated on over 13,000 thermal frames spanning diverse speeds, geometries and defect types, the system meets aerospace-grade benchmarks, eliminates robot-coupled dependencies, and delivers scalable, real-time AFP quality inspection directly from thermal imagery.

## ÖZET

# OTOMATİK FIBER YERLEŞTİRMEDE YERİNDE TERMAL DENETİM İÇİN ÖĞRENME TABANLI UÇTAN UCA HIBRİT BİR SİSTEM

MUHAMMED ZEMZEMOĞLU

Mekatronik Mühendisliği Doktora Tezi, Temmuz 2025

Tez Danışmanı: Prof. Dr. MUSTAFA ÜNEL

Anahtar Kelimeler: Otomatik Fiber Yerleştirme (AFP), Yerinde İzleme, Termografik Denetim, Kusur Tanımlama, Kalite Değerlendirme, Hareket Kestirimi, Serim Yeniden Yapılanma, Bilgisayarlı Görü, Derin Öğrenme

Otomatik Fiber Yerleştirme (AFP) teknolojisi, kompozit üretiminde öncü ve dönüştürücü bir rol üstlenmeye devam etmektedir. Ancak, süregelen kalite kontrol zorlukları bu gelişimi kısıtlamaktadır. Üretim sırasında ortaya çıkan malzeme ve imalat kaynaklı kusurlar, yapısal bütünlüğü tehdit edebilmekte; buna karşın mevcut denetim yöntemleri genellikle manuel, tepkisel ve gerçek zamanlı geri bildirim sağlamaktan uzaktır. Bu durum, üretim hatalarını tespitte gecikmelere ve yüksek maliyetli makine duruşlarına yol açmaktadır. Mevcut yöntemler iki temel eksiklik taşımaktadır: (1) yalnızca kısıtlı görevleri kapsayan, çoğunlukla çevrimidisi çalışan kısıtlı çözümler ve (2) zaman içinde evrilen katmanlar boyunca bütünsel ve tutarlı bir görselleştirme eksikliği. Bu doktora tezinde, yalnızca termal görüntüleme temelli, makineden bağımsız çalışan, gerçek zamanlı ve öğrenme odaklı çift yapılı bir denetim sistemi önerilmektedir. Sistem, akıllı yerel kusur tespiti ile küresel serim izlenebilirliğini bütünlük şekilde sunmaktadır.

Birinci yapı, üç bütünlük modülden oluşan, hibrit bir görüntü tabanlı kusur analiz ve kalite değerlendirme sistemi: Dinamik Şerit Tanımlama, Hiyerarşik Kusur Tanımlama ve Serim Kalitesi Değerlendirme. Bu modüller, çalışma zamanı (runtime) verimliliği için paralel veya koşullu çalışacak şekilde yapılandırılmıştır. Kullanımdan bağımsız özgün bir uzamsal-zamansal analiz algoritması sayesinde şerit sınırları piksel-altı hassasiyetle (ortalama hata  $< 0,8$  piksel) belirlenerek şerit düzeyinde

karar mekanizmaları etkinleştirilir. Aynı anda, üst seviyeli kusur tespiti, dokusal özelliklere dayalı optimize edilmiş bir Gabor-SVM sınıflandırıcı ile gerçekleştirilir (doğruluk ve duyarlılık  $> 95\%$ ). Kusurlu olarak işaretlenen görüntüler, özel olarak tasarlanmış 12 katmanlı derin evrişimsel sinir ağına (CNN) yönlendirilerek çok sınıflı alt seviye kusur sınıflandırması yapılır (doğruluk: 96,4%). Kusur tanımlama sonrası, termal uygulamaya uyarlanmış tohumlu aktif kontur algoritması ile segmentasyon gerçekleştirilir; elde edilen maskeler 93,2 piksel doğruluğu ile gürültülü ortamlarda dahi yorumlanabilir sonuçlar üretir. Bu çıktı, şerit geometrisi ile birleştirilerek, kusurun hem şerit hem de hat seviyesindeki etkisini sayısal olarak tanımlayan yeni bir metrik olan Kusur Etkilenim Oranı (DAP) hesaplanır ve bu metrik, operatör karar destek sisteminin temelini oluşturur. 5 fps hızla çalışan bu yapı, tam otomatik ve yorumlanabilir denetim olağanı sunarak havacılık standartlarını karşılamaktadır.

İkinci yapı olan ThermoRAFT-AFP, AFP sürecine özel olarak geliştirilmiş, makineden bağımsız çalışan, derin öğrenme tabanlı bir termal hareket kestirim birimidir. RAFT optik akış mimarisi üzerine inşa edilen bu yapı, AFP'ye özgü iyileştirmeler ve süreç-farkındalık optimizasyonlarla (örneğin öngörüye dayalı ilklendirme, sürüklendirme düzeltmesi, uyarlamalı erken çıkış) güçlendirilmiştir. Sistem, görüntüler-arası yoğun hareket tahminleri yaparak iki aşamalı bir yeniden yapılandırma hattını tetikler: önce serim şeridi düzeyinde termal mozaikler oluşturulur, ardından bunlar yüksek doğruluklu, katman seviyesinde serim görsellerine dönüştürülür. Bu yapı, zaman tutarlılığı ve global hizalamayı geri kazanarak kusur izlenebilirliğini mümkün kılar ve endüstriyel bakım-denetim mantığını yansıtır. ThermoRAFT-AFP, 4,83 mm/s hız tahmini ortalama kare hatası (RMSE) ve 0,1'in altında birikimli sürüklendirme ile, SNR = 14,4 dB seviyelerine kadar dayanıklılığını korurken 25 fps hızında gerçek zamanlı çalışmayı sürdürmektedir. Robotla senkronize verilere ya da düşük çözünürlüklü akışlara dayalı geleneksel yöntemlerin aksine, hiçbir yeniden eğitim gerektirmeden farklı üretim koşullarına başarıyla uyum sağlamaktadır.

Bu iki yapının birleşimi, AFP için gerçek zamanlı, bütünselik ve termal tabanlı bir denetim sistemi oluşturmaktadır. Yerel kusur çıktıları, global termal serim görselleri üzerine bindirilerek zaman uyumlu ve yorumlanabilir bir kalite haritası ortaya konmuştur. Böylece kusur tanımlama, segmentasyon ve hareket kestirimini birleşik bir izleme mimarisi içinde bütünlüğe getirmektedir. 13.000'den fazla termal görüntü üzerinde test edilen sistem, farklı hızlar, geometri ve kusur sınıfları altında havacılık düzeyinde doğruluk ve genellenebilirlik sergileyerek yeniden eğitime ihtiyaç duymadan çalışmaktadır; yalnızca termal görüntülerle, robot verisi olmaksızın, gelecek nesil AFP denetimi için ölçülebilir ve pratik bir çözüm sunmaktadır.

## ACKNOWLEDGEMENTS

This dissertation marks the end of a long academic journey—one that has shaped me both professionally and personally. Along the way, I encountered challenges, received invaluable support, and grew through experience. I am sincerely grateful to the many individuals and institutions whose guidance, encouragement, and contributions made this work possible.

I extend my deepest appreciation to my advisor, Prof. Dr. Mustafa Ünel, for his unwavering support, intellectual guidance, and generous mentorship. His distinctive lecturing style, technical rigor, and relentless commitment to excellence made him a lasting source of inspiration and an enduring role model. His impact transcended supervision—anchored in his faith in my potential and his persistent drive to raise my standards, challenge my limits, and guide me toward the best version of myself. It has been a privilege to pursue this research under his supervision.

I am sincerely grateful to my dissertation committee for their time, insightful feedback, and valuable contributions to the depth and quality of this work.

I gratefully acknowledge the support of TÜBİTAK under Grant No. 218M715, which made this research possible. I also thank Sabancı University for its intellectually rich and supportive environment, and my colleagues at SU-IMC and the project team for their dedication and collaboration. I am especially thankful to my friends in the Control, Vision, and Robotics (CVR) research group for their camaraderie and technical exchange. In particular, I thank my fellow Ph.D. companions—Naida Fetic, Mehmet Emin Mumcuoğlu, Shawqi Mohammed Farea, and Asmaa Loulou—for sharing this journey. And to my close friends whose presence made even the heaviest days lighter—thank you.

To my family, whose love, prayers, and quiet strength carried me through—your support has been my foundation. To my wife, Fatma—I lived my happiest days by your side. What we shared runs deeper than words, and no gift compares to the one you gave me: Yasemin—my little miracle, whose smile brightened the darkest nights and gave meaning to every effort.

As this chapter ends, I move forward with more than a dissertation—I carry a better version of myself, shaped by grace, grit, and gratitude.

— M.Z.

*To the miracle of my life... you, Yasemin...*

*To family... ever present, ever precious...*

*To the martyrs... light upon our path...*

## TABLE OF CONTENTS

<b>LIST OF TABLES .....</b>	<b>xiii</b>
<b>LIST OF FIGURES .....</b>	<b>xiv</b>
<b>LIST OF ABBREVIATONS .....</b>	<b>xvii</b>
<b>1. INTRODUCTION.....</b>	<b>1</b>
1.1. Absence of a Real-Time End-to-End Defect Inspection System .....	3
1.2. Machine Dependence and Lack of Global Motion Context .....	4
1.3. Bridging the Gap and Promoting Hybrid, Learning-Based Solutions..	5
1.4. Problem Formulation and Research Objectives .....	6
1.4.1. Framework 1: Frame-Wise Defect Analysis and Quality As- sessment .....	6
1.4.2. Framework 2: Independent Motion-Aware Global Inspection ..	7
1.4.3. Unified Research Objectives.....	8
1.5. Key Contributions .....	9
1.5.1. Framework 1: In-Situ Defect Analysis and Quality Assessment	9
1.5.2. Framework 2: Motion Estimation and Lay-Up Reconstruction	10
1.6. Thesis Outline and Organization .....	12
1.7. Publications.....	13
<b>2. LITERATURE REVIEW .....</b>	<b>14</b>
2.1. Sensor Modalities and Imaging Paradigms for AFP Inspection .....	14
2.2. Classical Thermographic Inspection and Its Limitations .....	15
2.3. Preprocessing Strategies: Raw vs. Rectified Perspectives .....	15
2.4. Learning-Based Defect Detection and Classification .....	16
2.5. Hybrid Strategies and Modular Pipelines.....	17
2.6. Temporal Reasoning and Motion-Coupled Inspection .....	17
2.7. Deep Optical Flow for Thermographic Sequences .....	18
2.8. Benchmarking and Dataset Availability .....	18
2.9. Synthesis and Research Gaps .....	19

<b>3. IN-SITU THERMOGRAPHIC MONITORING SETUP FOR AFP INSPECTION .....</b>	<b>20</b>
3.1. Thermographic Inspection in AFP Processes .....	20
3.2. In-Situ Thermal Vision Setup Design and Integration .....	21
3.2.1. Sensor Requirements .....	22
3.2.2. Flexible Holder Mechanism Design.....	22
3.2.3. Integration to AFP Machine Head .....	23
3.3. Closed-Loop Thermal Conditioning System .....	25
3.4. Thermal Camera Spatial Calibration .....	25
3.4.1. Correspondence Problem .....	26
3.4.2. Calibration Setup and Results .....	27
3.5. Experimental Setup and Data Acquisition Protocol .....	29
3.6. Synthetic Defect Scenarios .....	30
3.7. Universal Preprocessing Pipeline .....	31
<b>4. HYBRID FRAME-WISE DEFECT ANALYSIS AND QUALITY ASSESSMENT FRAMEWORK .....</b>	<b>35</b>
4.1. Framework Overview .....	35
4.2. Dynamic Tow Identification .....	38
4.2.1. Spatial-Temporal Analysis Algorithm .....	39
4.3. Hierarchical Defect Identification .....	41
4.3.1. High-Level Defect Detector .....	41
4.3.1.1. Gabor-Based Feature Extraction .....	42
4.3.1.2. SVM Classifier Model .....	43
4.3.2. Low-level Defect Classifier .....	45
4.3.2.1. Deep Learning Network Architecture.....	45
4.3.2.2. Model Training and Refinement .....	46
4.4. Lay-up Quality Evaluation.....	46
4.4.1. Defect Segmentation and Localization .....	47
4.4.2. Defect Area Percentage Estimation .....	48
<b>5. MACHINE-INDEPENDENT MOTION-AWARE GLOBAL INSPECTION FRAMEWORK.....</b>	<b>50</b>
5.1. Framework Overview .....	50
5.2. Deep Flow Estimation in AFP.....	52
5.3. ThermoRAFT-AFP: Domain Tailored Motion Estimator .....	54
5.3.1. Predictive Warm-Start Initialization .....	54
5.3.2. Regression-Based Drift Correction .....	54
5.3.3. Temporal Stabilization .....	55
5.4. Process-Aware Inference Optimizations .....	56

5.4.1. Region-of-Interest Cropping.....	56
5.4.2. Adaptive Early Exit .....	57
5.4.3. Temporal Subsampling .....	57
5.5. Reconstructing the AFP Process .....	58
5.5.1. Course-Level Seamless Image Stitching .....	58
5.5.2. Laminate-Level Assembly via Motion Synchronization .....	60
<b>6. EXPERIMENTAL RESULTS AND DISCUSSION .....</b>	<b>61</b>
6.1. Multi-Resolution Thermal Dataset for Dual Framework Evaluation...	61
6.1.1. Framework 1 Evaluation Dataset:.....	62
6.1.2. Framework 2 Evaluation Dataset:.....	62
6.2. Framework 1 Evaluation: Defect Analysis and Quality Assessment ...	62
6.2.1. Performance Analysis of Dynamic Tow Identification Algorithm	63
6.2.2. Hierarchical Defect Identification Performance Evaluation.....	66
6.2.2.1. High-Level SVM-Based Defect Detector.....	66
6.2.2.2. Low-Level DCNN-Based Multiclass Classifier .....	67
6.2.3. Defect Segmentation and Localization .....	69
6.3. Framework 2 Evaluation: Motion-Aware Global Inspection .....	71
6.3.1. ThermoRAFT-AFP Performance Analysis.....	72
6.3.2. Runtime Performance and Real-Time Feasibility .....	74
6.3.3. Lay-up Reconstruction and Visualization .....	76
6.3.4. Reliability and Robustness Analysis .....	78
6.3.5. Benchmarking of Motion Estimation Methods .....	79
6.4. System-Level Qualitative Comparison .....	81
<b>7. INTEGRATED SYSTEM OPERATION: PRACTICAL CASE STUDIES DEMONSTRATION .....</b>	<b>84</b>
7.1. Case Study A: Unidirectional Lay-up with Multiple FODs .....	85
7.2. Case Study B: Cross-Ply Lay-up with Mixed Defect Types .....	87
<b>8. CONCLUSIONS AND FUTURE DIRECTIONS .....</b>	<b>90</b>
<b>BIBLIOGRAPHY.....</b>	<b>93</b>

## LIST OF TABLES

Table 3.1. AFP Process and Environmental Conditions.....	30
Table 4.1. Optimized DCNN for low-level defect classification .....	45
Table 6.1. Mean performance comparison of the proposed spatial-temporal analysis based dynamic tow identification algorithm using different filtering methods and averaging techniques.....	65
Table 6.2. SVM model optimized parameters and performance evaluation.	67
Table 6.3. DCNN model parameters and classification results.....	68
Table 6.4. Performance metrics of ThermoRAFT-AFP across evaluation dataset .....	74
Table 6.5. Accuracy vs inference time trade-off analysis for different Step-Size settings.....	75
Table 6.6. Metric-wise comparison of ThermoRAFT-AFP performance across lay-up depth and feature richness. Shading reflects metric magnitude (lower is better). Arrows ( $\downarrow$ ) mark best values. $\Delta H-D$ is the healthy-to-defective delta. ....	79
Table 6.7. Metrology robustness of ThermoRAFT-AFP across noise levels ( $\sigma^2$ , SNR). Color shading qualitatively reflects performance drop per metric. The 14.4 dB row marks the final noise condition with acceptable estimation fidelity.....	80
Table 6.8. Comparison of motion estimation methods on thermal AFP dataset. Metrics span accuracy, runtime, and deployability. Bold marks best scores. ....	80
Table 6.9. System-level functional coverage comparison of AFP inspection systems across recent studies. Capabilities span defect analytics, thermal motion tracking, and real-time deployment. ....	82

## LIST OF FIGURES

Figure 1.1. (a) A typical AFP system during composite part production.	1
(b) A simplified diagram illustrating the working mechanism of the AFP robotic head while depositing a row of fiber tows. ....	1
Figure 1.2. An illustration of common AFP process defects. ....	2
Figure 1.3. The design–manufacture–inspect cycle for AFP processes.....	4
Figure 3.1. Thermal contrast patterns associated with common AFP defects.	21
Figure 3.2. Assembled thermal vision holder mechanism: (Left) open functional configuration, (Right) final with protective casing. ....	23
Figure 3.3. Manufactured thermal vision system without protective casing.	24
Figure 3.4. Integrated thermal vision system mounted on the AFP machine head.....	24
Figure 3.5. Tooling surface temperature map showing a $+7^{\circ}\text{C}$ edge–center gradient caused by imperfect contact with heating plates. ....	26
Figure 3.6. Thermal calibration mask: (a) CAD design, (b) waterjet cutting process, (c) manufactured composite mask.....	27
Figure 3.7. Thermal calibration setup: (a) Assembled rig with labeled components; (b) thermal image showing strong mask–background contrast at $150^{\circ}\text{C}$ . ....	28
Figure 3.8. Thermal calibration extrinsics: (a) camera-centered view; (b) world-centered view. ....	29
Figure 3.9. Flowchart of the universal preprocessing pipeline. ....	32
Figure 3.10. Temperature normalisation: (a) perspective corrected image; (b) temperature-normalised version showing deviation from row-wise mean. ....	33
Figure 3.11. (a) Raw thermal image of a flawless lay-up; (b) corresponding dataset-grade preprocessed result. ....	33
Figure 3.12. Examples of preprocessed defective lay-up instances: (a) gap (b) overlap (c) missing tow (d) tow splice and (e) foreign body. ....	34

Figure 4.1. Architecture of the proposed Hybrid Frame-wise Defect Analysis and Quality Assessment Framework, illustrated using a sample FOD defective thermal input.....	36
Figure 4.2. Tow boundary analysis in thermographic AFP inspection: (a) 3D thermal surface map with inter-tow valleys; (b) cross-sectional intensity profiles showing local minima aligned with tow boundaries ..	38
Figure 4.3. Gabor filter-based feature extraction: (a) $5 \times 8$ filter bank; (b) sample convolution outputs from a defective frame. ....	43
Figure 4.4. Soft-margin SVM decision boundary separating healthy and defective classes in feature space. ....	44
Figure 4.5. Preprocessing via mean subtraction: (a) input image with defect (FOD), (b) mean image of healthy lay-ups, (c) result after subtraction enhancing defect contrast.....	47
Figure 5.1. Overall structure of the proposed Motion-Aware Global Inspection Framework, illustrating the thermal input stream, motion estimation process, and reconstructed laminate output. ....	51
Figure 5.2. RAFT optical flow architecture.....	53
Figure 6.1. Tow boundary estimation using spatial-only analysis on (a) healthy and (b) defective samples with ground truth (green lines) and estimated boundaries (blue dashed lines); bar charts show edge distributions.....	63
Figure 6.2. Tow boundary errors under defects D1 and D2. Temporal averaging stabilizes B2 and B3 within $\pm 2$ px band.....	64
Figure 6.3. Normalized confusion matrix for SVM-based test set classification results. ....	67
Figure 6.4. Normalized confusion matrix for DCNN-based test set classification results.....	68
Figure 6.5. Adaptive thresholding of texture features: (a) original lay-up image showing a tow scrap foreign body defect; (b) texture feature-enhanced image; (c) segmented mask using Otsu's adaptive thresholding. ....	69
Figure 6.6. Active contours with texture feature seeding: (a) initial seed mask at $i = 1$ ; (b) intermediate contour at $i = 10$ ; (c) final convergence at $i = 25$ , isolating the defect. ....	70
Figure 6.7. Motion estimation for a representative AFP course: (top) frame-wise lay-up velocity, (bottom) cumulative displacement.....	73

Figure 6.8. Course-level reconstruction for a uni-directional laminate containing both healthy and FOD-affected regions. Stitched frame centers are marked by transparent dashed white lines for reference. Only a representative segment of the full course is shown for visual clarity and space efficiency.....	76
Figure 6.9. Ply-level motion synchronization visualization across two stitched AFP courses. Overlayed motion profiles (blue) highlight instantaneous velocity variations along the course length. Estimated inter-course lag is shown with a scaled marker in physical units (mm), derived from cross-correlation.....	77
Figure 6.10. Bland–Altman analysis of frame-wise displacement estimates, showing deviation mean (black), 95% limits (orange), and $\pm \frac{1}{2}$ pixel bounds (blue). Over 91% of estimates fall within the subpixel range; outliers coincide with velocity transitions.....	78
Figure 7.1. Case A: Unidirectional $[0^\circ]$ laminate with tow scrap FODs. Right: exploded-view reconstruction showing vertical thermal propagation. Left: zoom-in on ply 2 highlighting lateral thermal separation zones.....	86
Figure 7.2. Frame-wise inference on a defective instance from Case A. Thermal anomaly is segmented and localized with high tow-wise DAP. The decision system recommends process stoppage and defect removal.....	86
Figure 7.3. Case B: Cross-ply $[0^\circ, 90^\circ, 0^\circ, 90^\circ]$ laminate with mixed defect types. Exploded 3D view reveals (a) tow splice, (b) delamination trail, and (c) cut-induced gap and accompanying overlap, demonstrating system accuracy under directional variation.....	88
Figure 7.4. Frame-wise inference on a defective instance from Case B. Localized tow gap triggers DAP-based alert and recommends corrective action via the knowledge base.....	89

## LIST OF ABBREVIATONS

<b>AC</b> Active Contours .....	47, 69, 70, 71
<b>AFP</b> Automated Fiber Placement .	xi, xiii, xiv, xv, 1, 9, 20, 21, 22, 23, 24, 29, 30, 31, 33, 35, 38, 41, 42, 48, 49, 50, 52, 53, 54, 55, 56, 57, 58, 63, 67, 69, 71, 76, 78, 79, 80, 81, 84, 85, 87, 89
<b>AFP-DSS</b> Automated Fiber Placement process decision-support system..	37, 49, 85, 87, 91
<b>AT</b> Adaptive Thresholding.....	47, 69, 70, 71
<b>C<sub>avg</sub></b> Cumulative Average.....	40
<b>CFRP</b> Carbon Fiber-Reinforced Plastic .....	1
<b>CNN</b> Convolutional Neural Network.....	45, 46
<b>CPU</b> Central Processing Unit .....	29
<b>DAP</b> Defect Area Percentage	xvi, 37, 48, 49, 64, 69, 70, 71, 84, 85, 86, 87, 89, 90, 91, 92
<b>DCNN</b> Deep Convolutional Neural Network .....	xiii, xv, 45, 66, 67, 68, 85
<b>FOD</b> Foreign Object Debris .....	xii, xv, xvi, 6, 31, 33, 36, 47, 76, 84, 85, 86
<b>FPS</b> Frames Per Second.....	71, 80, 81
<b>GF</b> Gabor Filter .....	42, 43, 44, 47, 48
<b>GPU</b> Graphical Processing Unit .....	29, 53
<b>GUI</b> Graphical User Interface .....	84
<b>IoU</b> Intersection over Union.....	70, 71
<b>ML</b> Machine Learning .....	41

<b>NDT</b> Non-Destructive Testing.....	20
<b>PA</b> Pixel Accuracy.....	70, 71
<b>RAFT</b> Recurrent All-Pairs Field Transforms.....	xv, 52, 53, 54, 56
<b>RBF</b> Radial Basis Function.....	44
<b>ReLU</b> Rectified Linear Unit .....	45
<b>RMSE</b> Root Mean Square Error.....	74, 75
<b>RoI</b> Region of Interest .....	22, 32, 51, 56, 58
<b>SGDM</b> Stochastic Gradient Descent with Momentum .....	46
<b>SM<sub>avg</sub></b> Simple Moving Average .....	40
<b>SVM</b> Support Vector Machine.....	xiii, xv, 43, 44, 66, 67, 85
<b>WM<sub>avg</sub></b> Weighted Moving Average.....	40

## 1. INTRODUCTION

AFP has become a cornerstone of high-performance composite manufacturing. Recent studies highlight that AFP technology accounts for approximately 50% of all aerospace composite construction (Oromiehie, Prusty, Compston & Rajan (2019)), underscoring its pivotal industrial impact. By combining robotic precision with advanced materials, AFP enables the layer-by-layer deposition of CFRP epoxy strips—known as *tows*—to fabricate lightweight, structurally optimized components with micron-level control (August, Ostrander, Michasiow & Hauber (2014)). These tows are laid in rows to form a *course*; multiple courses build a *ply*, and stacked plies collectively form the laminate. The AFP head typically integrates a compaction roller and a heat source to ensure bonding and adhesion during each pass, forming a tightly controlled additive manufacturing sequence, as illustrated in Figure 1.1b, adapted from (Juarez, Gregory & Cramer (2018)).

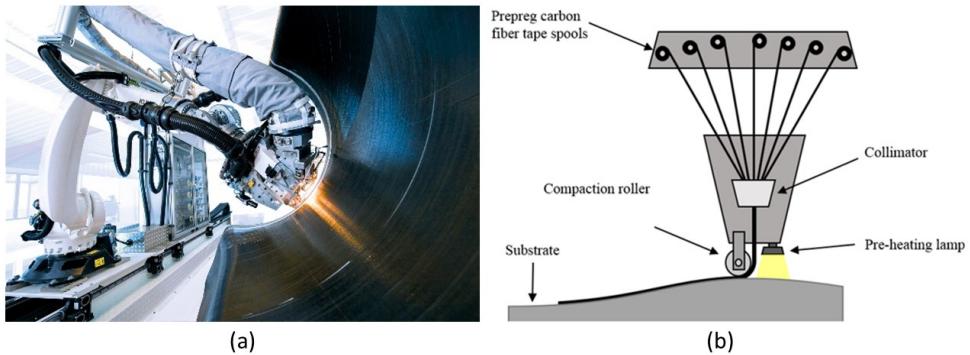


Figure 1.1 (a) A typical AFP system during composite part production. (b) A simplified diagram illustrating the working mechanism of the AFP robotic head while depositing a row of fiber tows.

AFP offers unmatched advantages in design freedom, production throughput, and structural performance—particularly in aerospace, automotive, and renewable energy sectors (Carosella, Hügle, Helber & Middendorf (2024))—but also introduces significant process control challenges (Sun, Han, Fu, Jin, Dhupia & Wang (2020)). Even minor deviations in compaction, heating, or toolpath execution can lead to critical manufacturing defects such as gaps, overlaps, wrinkles, misalignments, and

embedded foreign objects (Anay, Miller, Tessema, Wehbe, Ziehl, Tatting, Gurdal, Harik & Kidane (2022)). These are illustrated in Figure 1.2, adapted from (Denkena, Schmidt, Völtzer & Hocke (2016)), and are widely reported to compromise interlaminar strength, reduce fatigue life, and—under severe conditions—degrade structural integrity by more than 30% (Abouhamzeh, Nardi, Leonard & Sinke (2018); Blom, Lopes, Kromwijk, Gurdal & Camanho (2009)).

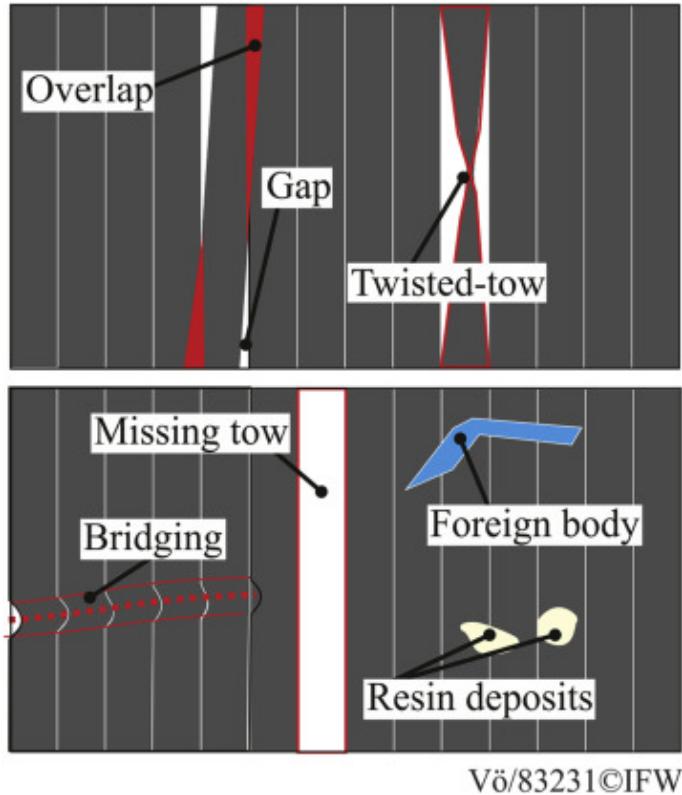


Figure 1.2 An illustration of common AFP process defects.

Although deposition is highly automated, quality control in AFP remains predominantly manual (Rudberg, Nielson, Henscheid & Cemenska (2014)). Operators are required to pause the process frequently to visually inspect each layer for defects. Although critical for quality assurance, this approach is time-consuming, error-prone, and incurs high labor costs (Cemenska, Rudberg & Henscheid (2015)). Recent surveys report that over 42% of AFP build time is devoted to inspection-related tasks—more than twice the time consumed by the actual lay-up process (Rudberg, Cemenska & Sherrard (2019)). Delayed or missed defect identification further exacerbates downstream repair complexity and jeopardizes structural compliance.

As AFP continues to scale toward more complex geometries and larger lay-up surfaces, maintaining consistent in-situ quality becomes increasingly challenging—not only for defect detection, but also for timely, interpretable feedback that supports production decisions. These limitations underscore the urgent need for an auto-

mated, in-situ monitoring solution capable of delivering real-time quality insights, minimizing downtime, enhancing part reliability, and directly supporting operator decision-making—aligned with Industry 4.0 paradigms of intelligent, data-driven manufacturing (Parmar, Khan, Tucci, Umer & Carlone (2022)). This brings forth a fundamental challenge at the intersection of automation, reliability, and scalability:

**How can we empower AFP systems with real-time, end-to-end inspection capabilities that are robust, intelligent, and seamlessly integrable—without sacrificing speed, accuracy, or generalizability?**

In this context, several critical research gaps and practical challenges emerge:

### 1.1 Absence of a Real-Time End-to-End Defect Inspection System

Despite recent academic advances, existing AFP inspection methods remain fragmented and limited in scope. Most focus on partial stages—such as defect detection or classification—and often operate post-process, require manual oversight, or depend on operator expertise. These approaches lack real-time responsiveness, offer no unified pipeline, and are rarely scalable across diverse setups.

As a result, the broader *design-manufacture-inspect* cycle remains disconnected, as illustrated in Figure 1.3. Feedback mechanisms are manual, optimization loops are delayed, and inspection rarely informs upstream design or process control decisions.

No commercially available or reported solution currently delivers a unified, thermographic, real-time inspection system that can identify defects as they emerge, quantify their severity, and guide operators using interpretable visual overlays and actionable frame-wise metrics (Yadav & Schledjewski (2023)). Closing this gap calls for a fully automated, in-situ, and learning-enabled pipeline capable of supporting both defect-level intelligence and process-level traceability.

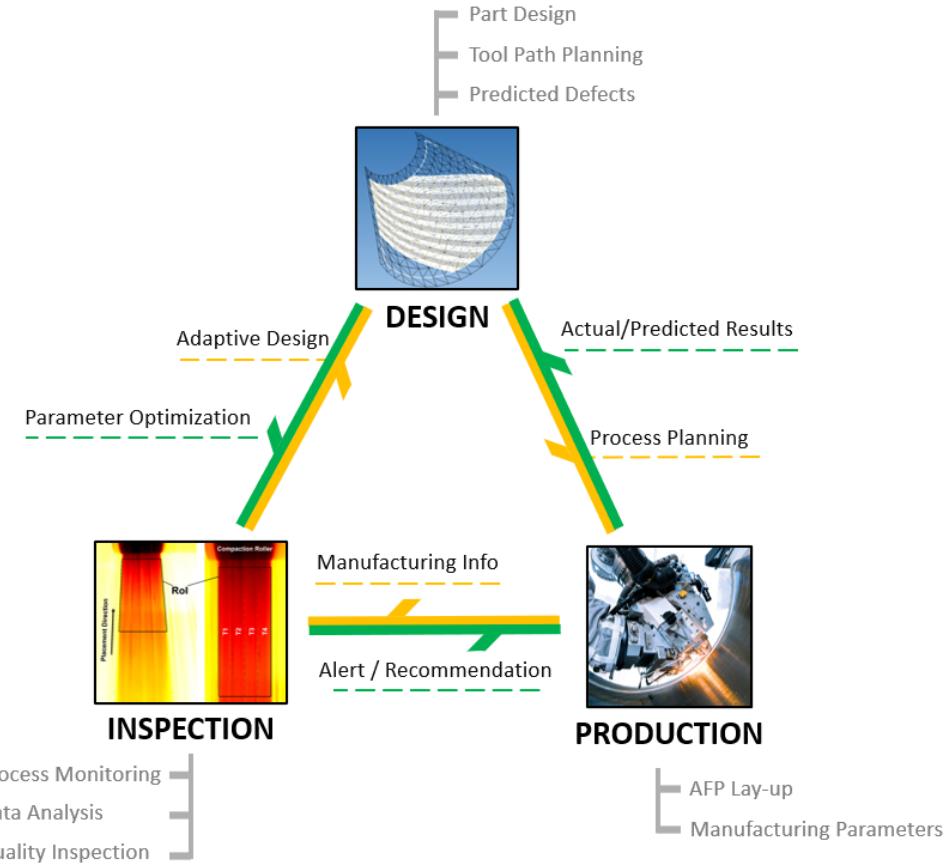


Figure 1.3 The design–manufacture–inspect cycle for AFP processes.

## 1.2 Machine Dependence and Lack of Global Motion Context

Localized, frame-wise inspection offers valuable defect-level insights but lacks the spatial and temporal continuity required for comprehensive quality assurance. Many defects evolve gradually, recur intermittently, or reveal structural trends only when viewed across time. Without motion-aware inspection, such phenomena remain undetected, hindering traceability, limiting contextual interpretation, and weakening the feedback loop to operators and process controllers.

Although some studies have introduced global lay-up visualization through motion information, they often depend on external signals—such as robot controller data or toolpath encoders—which introduce synchronization errors and restrict generalizability to specific AFP machines (Hocke (2020); Juarez & Gregory (2021)). Other efforts based on thermal optical flow—such as those relying on handcrafted curvature descriptors—remain limited in adaptability and robustness when applied across varying AFP setups and defect scenarios (Denkena, Schmidt, Timmermann & Friedel (2022)).

At present, no deployable, machine-independent approach exists for estimating motion directly from thermal image sequences to support global inspection. Bridging this gap would enable spatially consistent lay-up visualization, defect traceability across time, and intuitive overlays that assist operators in interpreting evolving laminate structures—laying the foundation for interpretable, scalable, and fully automated AFP quality monitoring.

### 1.3 Bridging the Gap and Promoting Hybrid, Learning-Based Solutions

To overcome the brittleness and low generalizability of traditional rule-based methods, learning-based solutions are increasingly favored for AFP quality inspection (Brasington, Sacco, Halbritter, Wehbe & Harik (2021)). These approaches can adapt to defect variability and thermal inconsistencies—but their effectiveness has been historically limited by the scarcity of representative thermal datasets.

In response, this work leverages an extensive in-house thermal dataset—captured under real AFP conditions—to enable a hybrid inspection strategy. By combining traditional computer vision, machine learning, and deep learning in a task-optimized architecture, the system balances interpretability, speed, and adaptability, making it suitable for industrial deployment.

Crucially, the proposed solution goes beyond frame-level analysis. It integrates localized defect insights with global ply visualization—providing spatially consistent overlays and operator-assistive feedback in real-time. This integration paves the way toward a new class of thermographic inspection systems that are learning-centric, modular, and fully automated.

## 1.4 Problem Formulation and Research Objectives

Building upon the challenges highlighted in the motivation section, this thesis addresses two critical gaps in the domain of thermographic inspection for Automated Fiber Placement (AFP). These challenges are tackled via a parallel dual-framework architecture, in which each framework is designed to address a major research gap and a fundamental aspect of the inspection problem: frame-wise, real-time defect analysis, and motion-aware, machine-independent global inspection. Together, these frameworks form a unified, operator-assistive, learning-centric quality monitoring system.

### 1.4.1 Framework 1: Frame-Wise Defect Analysis and Quality Assessment

The first challenge centers on enabling reliable, real-time defect analysis directly from thermal images acquired during the AFP process. Defects such as gaps, overlaps, missing tows, and FOD can appear at any time and must be identified promptly with sufficient spatial precision and semantic clarity to support subsequent quality assessment. The research problem addressed in this framework is formulated as:

*How can thermal frames acquired during AFP be processed in real time to detect, classify, and localize multiple defect types—while ensuring tow-level interpretability, pixel-level accuracy, and industrial deployability without human intervention?*

To address this problem, Framework 1 is formulated as a hybrid, frame-wise inspection pipeline. It combines classical computer vision, machine learning, and deep learning to progressively extract semantic and spatially detailed information from each frame. Its outputs form the local intelligence layer of the proposed system, offering immediate qualitative and quantitative insights into defect types and spatial locations. To fulfill this vision, the framework is designed around the following core research objectives:

- **Robust Tow Identification:** Accurately and consistently estimate tow boundaries in noisy and low-contrast thermal images, without requiring prior

assumptions or manual tuning.

- **Reliable Defect Detection and Classification:** Detect defective lay-ups across varied scenarios with high accuracy, prioritizing sensitivity to false negatives to ensure all defects are captured, and reducing false positives to avoid unnecessary downtime. Distinguish between multiple defect types with high precision to enable actionable, granular quality assessment.
- **Precise Segmentation:** Spatially localize each defect with sufficient resolution to support quantitative evaluation and facilitate targeted repair or rejection.
- **Real-Time Operation:** Ensure the framework achieves its objectives while operating within strict inference time constraints on standard industrial computing platforms.

#### 1.4.2 Framework 2: Independent Motion-Aware Global Inspection

The second challenge addresses the lack of temporal continuity and global spatial context in current thermographic inspection pipelines. While frame-wise analysis captures local defect phenomena, it fails to trace defects across time, detect full-length course-level misalignments, or visualize evolving laminate structures. This limitation is especially critical when parts are large, geometrically complex, or require cumulative surface reconstruction for quality traceability. The research problem tackled in this framework is defined as:

*How can motion in automated fiber placement be accurately and robustly estimated, and the lay-up reconstructed, using only thermal video data—without relying on machine signals—while preserving spatial coherence, temporal consistency, and enabling global reconstruction?*

To tackle this challenge, Framework 2 is designed as a motion-aware global inspection system based on thermal-domain optical flow. Unlike prior methods that depend on robot encoder data or process synchronization, this framework estimates displacement directly from overlapping thermal images. It reconstructs the lay-up surface in real time and anchors local defect information within a coherent global view. To achieve this, the framework pursues the following core research objectives:

- **Accurate Motion Estimation:** Estimate vertical displacement across thermal frame sequences with high pixel-level accuracy and temporal stability—even in low-texture, thermally inconsistent conditions.
- **Real-Time Global Reconstruction:** Process high-frequency thermal streams with minimal latency to enable uninterrupted ply visualization during the AFP lay-up.
- **Machine-Independent Deployment:** Operate robustly across different AFP setups and lay-up conditions without relying on kinematic inputs, robot-specific parameters, or external synchronization.
- **Spatially Coherent Visualization:** Facilitate ply-level defect traceability through interpretable, top-down thermal mosaics—enabling spatially anchored overlays and operator-assistive navigation.

#### 1.4.3 Unified Research Objectives

This dissertation proposes a unified, learning-centric inspection system that integrates localized frame-wise defect analysis with global thermographic lay-up reconstruction. The overarching research objectives are:

- **Enable Integrated Operation:** Design a dual-framework architecture that combines real-time defect detection and global laminate visualization into a cohesive, operator-assistive pipeline.
- **Ensure Interpretability and Real-Time Feedback:** Deliver spatially and temporally consistent inspection outputs that support intuitive monitoring, decision-making, and traceability.
- **Demonstrate Robustness Through Empirical Validation:** Evaluate the full system under diverse AFP conditions using a high-resolution thermal dataset comprising over 13,000 in-situ frames.

## 1.5 Key Contributions

This thesis makes foundational contributions to the field of automated composite manufacturing by introducing the first fully integrated, real-time, learning-centric thermographic inspection system for AFP processes. Built upon a dual-framework architecture, the proposed system redefines both local defect intelligence and global lay-up traceability—achieving full automation, machine-independence, and real-time performance from raw thermal imagery to operator feedback. The contributions are presented below in two categories aligned with the system’s core components.

### 1.5.1 Framework 1: In-Situ Defect Analysis and Quality Assessment

- **A Hybrid, Frame-Wise Defect Analysis System.** We present a dual-stage system that fuses classical computer vision with machine learning and deep learning to enable real-time, tow-level defect analysis from thermal imagery—delivering both semantic and spatial detail without manual intervention. This represents the first comprehensive learning-based thermal inspection pipeline optimized for defect intelligence in AFP.
- **Tow Boundary Estimation via Setup-Independent Spatial–Temporal Analysis.** A novel algorithm exploits the sequential nature of AFP deposition to dynamically identify tow boundaries with sub-pixel accuracy across evolving thermal frames. Free from rigid geometric assumptions, this method achieved a mean error of **0.8 px** with stable convergence.
- **Hybrid Hierarchical Defect Identification under Data Scarcity.** To address data imbalance and resource constraints, we design a hybrid classification pipeline that combines an SVM-based high-level defect detector (using Gabor texture features) with a 12-layer DCNN for fine-grained classification. Trained on an enriched dataset with synthetic augmentation, the system achieved **96.4% accuracy** across all classes.
- **Active Contour-Based Defect Segmentation for Spatial Precision.** We implement seeded active contour models specifically adapted to thermal image textures, achieving precise segmentation of multiple defect types. The method yields a mean **pixel accuracy of 93.2%** and **IoU of 0.72**, enabling

robust quality metric computation and spatially anchored visualization.

- **Defect Area Percentage (DAP) Metric and the AFP-DSS Operator Support Engine.** We introduce the novel **DAP metric** as a quantitative measure of defect severity at both course and tow levels. Embedded within an expert-driven decision-support system (**AFP-DSS**), this module provides real-time alerts and intelligent recommendations—enhancing process transparency and empowering automated operator assistance.

### 1.5.2 Framework 2: Motion Estimation and Lay-Up Reconstruction

- **ThermoRAFT-AFP: A Deep Learning-Based Thermal Motion Estimation System.** We introduce **ThermoRAFT-AFP**, the first machine-independent, learning-based solution for dense motion estimation and lay-up reconstruction in AFP using thermal imagery. Built around a thermally adapted RAFT backbone, it is specifically tailored to the kinematic, radiometric, and temporal characteristics of the AFP process using multiple customizations including predictive warm-start initialization, regression-based drift correction, exponential moving average (EMA) filtering.
- **Real-Time Operation Through Custom Process-Aware Optimization.** ThermoRAFT-AFP achieves real-time performance at **25 Hz** through a suite of custom strategies, including RoI gating, adaptive early exit, temporal subsampling, and hyperparameter auto-tuning. These ensure both stability and responsiveness on standard industrial hardware.
- **Laminate-Scale Reconstruction via Velocity-Aware Stitching.** We propose a novel reconstruction pipeline that combines vertical motion estimates with course-level stitching and velocity-aware alignment, producing continuous 2D thermal mosaics and interpretable 3D ply-wise visualizations—enabling full-laminate traceability and enhancing defect propagation analysis.
- **State-of-the-Art Accuracy and Robustness on Large-Scale Evaluation.** Validated on a dataset of over **13,300** in-house thermal frames, ThermoRAFT-AFP achieves **RMSE  $\approx 4.83 \text{ mm/s}$** ,  **$R^2 > 99.4\%$** , and  **$\text{MPE} \approx -0.67\%$** , while remaining resilient down to  **$\text{SNR} = 14.4 \text{ dB}$** . It consistently outperforms classical and AFP-specific methods across accuracy, robustness, and deployment-readiness.

## System-Level Contributions

- **Largest Annotated Thermal Dataset for AFP Inspection.** We construct the most extensive AFP-specific thermal dataset to date, comprising over 13,000 annotated frames captured under real production conditions. Task-specific synthetic defect scenarios were incorporated to rigorously support both defect-level analysis (Framework 1) and motion-based reconstruction (Framework 2).
- **First End-to-End, Real-Time Inspection Pipeline from Image to Operator.** This dissertation presents the first fully integrated system that transforms raw thermal input into actionable quality feedback—combining real-time defect intelligence with global laminate visualization and delivering operator-assistive overlays validated through realistic industrial case studies.
- **Empirical Validation across Accuracy, Runtime, and Deployability Metrics.** Both frameworks are quantitatively benchmarked for precision, execution time, and robustness across varied AFP conditions. The system consistently satisfies industrial-grade real-time constraints while maintaining high interpretability and generalization.

## 1.6 Thesis Outline and Organization

The remainder of this thesis is structured as follows:

**Chapter 2** reviews the current state of thermographic inspection for Automated Fiber Placement, covering conventional and learning-based methods for defect detection, classification, segmentation, and motion estimation, and identifying key gaps addressed in this thesis.

**Chapter 3** describes the in-situ thermographic monitoring platform developed for this study. It details the hardware setup, thermal camera calibration, synthetic defect scenario design, and the universal preprocessing pipeline that ensures standardized input across both proposed frameworks.

**Chapter 4** introduces the first core framework for frame-wise defect analysis and quality assessment. It outlines the hybrid design and learning-centric architecture that enable real-time, tow-level inspection from thermal imagery.

**Chapter 5** presents the second core framework for machine-independent lay-up reconstruction via thermal motion estimation. It details the ThermoRAFT-AFP model and its integration into a real-time pipeline for global laminate visualization.

**Chapter 6** provides a detailed evaluation of both frameworks, reporting quantitative and qualitative results across accuracy, runtime, robustness, and deployability. The evaluations are based on a high-resolution, intensive thermal dataset acquired under real AFP conditions.

**Chapter 7** demonstrates the real-time integrated operation of the complete system through two representative case studies. It illustrates how local frame-wise intelligence and global laminate reconstruction converge to produce actionable, operator-assistive visualizations.

**Chapter 8** concludes the thesis by summarizing key contributions and findings, and outlines future directions toward scalable, intelligent, and fully autonomous thermographic inspection systems for composite manufacturing.

## 1.7 Publications

The work presented in this thesis has contributed to the following peer-reviewed publications:

- **M. Zemzemoglu**, M. Unel, and L. T. Tunc, *Enhancing automated fiber placement process monitoring and quality inspection: a hybrid thermal vision based framework*, Composites Part B: Engineering, vol. 285, pp. 111753, 2024.
- **M. Zemzemoglu**, M. Unel, *Deep learning-based thermal motion estimation and lay-up reconstruction framework towards machine-independent real-time AFP process monitoring and inspection*, (Under Review), 2025.
- **M. Zemzemoglu**, M. Unel, *Thermal inspection and quality assessment for AFP processes via automatic defect detection and segmentation*, in Proceedings of the IEEE 50th Annual Conference of the Industrial Electronics Society (IECON), pp. 1–6, 2024.
- **M. Zemzemoglu**, M. Unel, *A hierarchical learning-based approach for the automatic defect detection and classification of AFP process using thermography*, in Proceedings of the IEEE 49th Annual Conference of the Industrial Electronics Society (IECON), pp. 1–6, 2023.
- **M. Zemzemoglu**, M. Unel, *Design and implementation of a vision-based in-situ defect detection system of automated fiber placement process*, in Proceedings of the IEEE 20th International Conference on Industrial Informatics (INDIN), pp. 393–398, 2022.

## 2. LITERATURE REVIEW

### 2.1 Sensor Modalities and Imaging Paradigms for AFP Inspection

Automated Fiber Placement (AFP) has emerged as a critical technology for manufacturing advanced composite structures with high geometric fidelity and production throughput. Despite these advantages, AFP remains vulnerable to process-induced defects—such as gaps, overlaps, bridging, and foreign object debris (FOD)—which necessitate robust in-situ monitoring systems for real-time quality assurance. A broad spectrum of sensing modalities has been proposed to this end, including ultrasonic probes, strain gauges, laser profilometers, and optical sensors (Drinkwater & Wilcox (2006); Oromiehie, Prusty, Rajan & Compston (2016); Schmidt, Denkena, Völtzer & Hocke (2017); Shadmehri, Ioachim, Pahud, Brunel, Landry, Hoa & Hojjati (2015)). Among these, laser profilometry has gained maturity for capturing surface anomalies (Koptelov, El Said & Tretiak (2025); Meister, Wermes, Stüve & Groves (2020)), but remains blind to thermally induced bonding issues and subsurface defects.

In contrast, vision-based systems—particularly those employing infrared thermography—have demonstrated promising versatility and non-contact operability. Thermographic inspection exploits the intrinsic heating phase in AFP to highlight thermal discontinuities that correlate with lay-up irregularities (Denkena, Schmidt & Weber (2016); Gregory & Juarez (2018); Juarez, Cramer & Seebo (2016)). Unlike visible-light imaging, which suffers from poor contrast on carbon fiber-reinforced polymer (CFRP) surfaces, thermal cameras yield richer spatiotemporal signatures that improve defect detectability and robustness across varying lighting and material conditions.

## 2.2 Classical Thermographic Inspection and Its Limitations

Early explorations of thermographic AFP inspection relied on handcrafted image processing techniques such as edge detection, filtering, and thresholding (Denkena et al. (2016); Juarez & Gregory (2019); Schmidt et al. (2017)). While computationally efficient, these methods demonstrated limited adaptability across process setups and defect types. Juarez and Gregory introduced high-pass and low-pass filter combinations to enhance operator visualization, but their system was not designed for automated detection (Juarez & Gregory (2019)). Similarly, Schmidt et al., proposed an edge-based segmentation pipeline, yet its heavy dependence on process-specific parameters rendered it brittle and unscalable (Schmidt et al. (2017)).

Critically, these classical methods failed to offer generalization across tow geometries, material types, or viewing angles—traits essential for industrial deployment. As AFP systems increase in complexity, inspection pipelines must transcend static rules and adopt adaptive, data-driven models capable of encoding thermal semantics and spatial priors.

## 2.3 Preprocessing Strategies: Raw vs. Rectified Perspectives

A foundational choice in thermal inspection systems concerns whether to analyze raw skewed frames or to apply geometric rectification. The first strategy processes thermal images as-is, avoiding calibration and preserving sensor integrity (Denkena et al. (2016); Schmidt et al. (2017); Yipeng, Wang, Wang, Li & Ke (2021)). However, it incurs geometric distortion that complicates downstream tasks like segmentation and defect localization. These pipelines often rely on threshold tuning specific to camera angles or fiber orientation, undermining robustness.

An alternative is perspective correction, often via homography, which reprojects the lay-up surface to a standardized, top-down view (Schmidt, Hocke & Denkena (2019)). This canonical alignment simplifies motion estimation and tow tracing by reducing spatial skew. Yet, this approach assumes fixed tow geometry and stable camera calibration—conditions often violated due to mechanical drift or toolpath variation. The trade-off between raw-frame fidelity and rectified-frame consistency

underscores the need for preprocessing pipelines that balance setup independence with geometric stability.

## 2.4 Learning-Based Defect Detection and Classification

The limitations of classical methods have led to a shift toward machine learning (ML) and deep learning (DL) techniques, which offer improved generalization across defect morphologies and process conditions. Classical ML systems have used texture features—such as Gabor filters—combined with support vector machines (SVMs) for binary classification. Deep convolutional neural networks (CNNs), by contrast, have demonstrated superior performance, with Schmidt et al. reporting over 90% classification accuracy on thermographic inputs for tow-level and course-level anomalies (Schmidt et al. (2019)).

Sacco et al. (Sacco, Baz Radwan, Anderson, Harik & Gregory (2020); Sacco, Radwan, Beatty & Harik (2019); Sacco, Radwan, Harik & Van Tooren (2018)) modified ResNet architectures to perform semantic segmentation on profilometry scans, attaining promising performance at the cost of high computational complexity. These results highlight a broader challenge: real-world deployment is constrained not by algorithmic accuracy alone but by runtime feasibility and dataset diversity. Most existing models are trained on small, institution-specific datasets, which hampers reproducibility and transferability.

To overcome data scarcity, some researchers have explored synthetic generation. Zambal et al. used U-Net models on synthetic depth maps (Zambal, Heindl, Eitzinger & Scharinger (2019)), while Meister et al. leveraged GANs for generating artificial data (Meister, Möller, Stüve & Groves (2021); Meister & Wermes (2023)). Despite improvements in training diversity, these models exhibit domain shift when applied to real thermographic inputs. Juarez et al., advocated for controlled defect injection in manufacturing trials to generate realistic yet reproducible training data—offering a pragmatic middle ground between realism and ground truth availability (Juarez & Gregory (2021)).

## 2.5 Hybrid Strategies and Modular Pipelines

Given the multifaceted nature of AFP defects, modular hybrid pipelines have been proposed to blend classical and learning-based approaches. Tang et al., proposed a two-step approach that uses geometry-based preprocessing followed by deep segmentation of 3D profilometry (Tang, Wang, Cheng, Li & Ke (2022)). Ghamisi et al., combined rule-based tow localization with unsupervised defect detection via neural autoencoders. These architectures show that modularity can enhance flexibility and interpretability (Ghamisi, Charter, Ji, Rivard, Lund & Najjaran (2023)).

Nonetheless, most hybrid methods lack temporal awareness and operate on static, frame-level inputs. Their components, though effective in isolation, are rarely integrated into cohesive systems with real-time guarantees, limiting their industrial applicability. This disjunction between task-level accuracy and pipeline-level deployability defines a critical gap in the current literature.

## 2.6 Temporal Reasoning and Motion-Coupled Inspection

Temporal context is essential for detecting defect progression, reconstructing ply maps, and improving localization accuracy. Several motion-aware systems have emerged. Juarez and Gregory synchronized thermal frames with machine control signals to reconstruct ply-scale views (Juarez & Gregory (2021)), while Hocke used controller feedback to anchor thermal defects in spatial context (Hocke (2020)). However, these approaches rely on precise synchronization and access to machine-side kinematic data, making them fragile and platform-dependent.

Denkena et al., proposed a machine-independent approach using handcrafted optical flow features based on local thermal curvature (Denkena et al. (2022)). While innovative, their system lacked real-time feasibility analysis, was validated only on a narrow proof-of-concept dataset, and did not address downstream integration with reconstruction or decision support. These limitations emphasize the need for scalable, self-contained systems that infer motion directly from thermal imagery without auxiliary machine feedback.

## 2.7 Deep Optical Flow for Thermographic Sequences

Recent progress in deep optical flow estimation offers new tools for thermographic motion tracking. Models such as FlowNet 2.0 (Ilg, Mayer, Saikia, Keuper, Dosovitskiy & Brox (2017)), PWC-Net (Sun, Yang, Liu & Kautz (2018)), and RAFT (Teed & Deng (2020)) have set performance benchmarks on visible-spectrum datasets, delivering dense, subpixel displacement fields through multi-level feature extraction and refinement (Wang, Wang, Li, Guo, Xu, Ma, Ling, Fu & Jia (2024)). These architectures have demonstrated robustness in texture-poor or occlusion-prone environments—traits that are theoretically transferable to thermal imaging.

Yet, direct application to thermography remains problematic. Thermal frames suffer from low contrast, emissivity variations, and high noise—features underrepresented in visible-spectrum training datasets. Additionally, models like RAFT rely on memory-intensive cost volumes and deep stacks of convolutional layers, resulting in execution times incompatible with real-time inspection (Alfarano, Maiano, Papa & Amerini (2024a)). No existing solution satisfies the trifecta of thermal specificity, subpixel accuracy, and runtime efficiency. These limitations motivate the development of domain-adapted, lightweight optical flow frameworks for thermal AFP inspection.

## 2.8 Benchmarking and Dataset Availability

Reproducibility and benchmarking are cornerstones of scientific progress. However, to date, there exists no publicly available thermographic dataset for AFP inspection. Proprietary restrictions and industrial confidentiality severely limit data sharing. Consequently, most research groups develop private, setup-specific datasets, which hinder comparison and generalization.

A partial exception is the multi-modal AFP dataset by Pantoji et al., which includes laser tracker, visible camera, and line laser sensor (LLS) data (Pantoji, Kassapoglou & Peeters (2025)). While valuable, this dataset excludes thermal imagery and thus cannot be used for thermography-based algorithm validation. In response to this deficit, we constructed a dedicated thermal AFP dataset under controlled labora-

tory conditions. This dataset supports reproducible evaluation of both frame-wise and motion-aware modules, and represents a critical resource for advancing thermo-graphic AFP inspection.

## 2.9 Synthesis and Research Gaps

The literature presents a fragmented landscape: individual advances in defect detection, motion estimation, or visualization exist, but are rarely integrated into deployable, real-time systems. Common bottlenecks include reliance on handcrafted features, limited dataset diversity, absence of temporal continuity, and machine-dependent architectures.

This dissertation addresses these gaps through a dual-framework architecture. The first stream performs hierarchical frame-wise defect analysis using adaptive, learning-based methods, while the second reconstructs global lay-up views via domain-tailored deep optical flow. Together, these modules operate in real-time, require no external machine data, and deliver a complete, interpretable thermal inspection system for AFP processes—laying the foundation for end-to-end intelligent quality assurance in composite manufacturing.

### 3. IN-SITU THERMOGRAPHIC MONITORING SETUP FOR AFP INSPECTION

Thermographic inspection underpins the proposed learning-centric system for AFP process inspection and quality assessment. This chapter presents the design, integration, and calibration of a fully in-situ thermal vision system that captures high-resolution, temporally rich imagery directly from the lay-up region. Operating independently of the AFP controller, the setup enables real-time observation of thermal signatures of the lay-up. These signatures drive the dual-framework architecture developed in this work: spatially standardized inputs drive local defect detection, classification, and segmentation, while temporally consistent sequences enable motion estimation and lay-up reconstruction. Together with environmental conditioning, camera calibration, and a universal preprocessing pipeline, this setup forms the foundation for all downstream inspection algorithms. Parts of this chapter, including figures and system descriptions, are adapted from the author’s earlier publication (Zemzemoglu & Unel (2022)).

#### 3.1 Thermographic Inspection in AFP Processes

Thermographic analysis is a NDT method that interprets surface heat distribution across a material. In AFP, it enables anomaly detection by capturing deviations in heat flux from the thermally conditioned tooling surface. Although the thermal field is driven by an external heat source, its observed pattern is shaped by factors such as deposition speed, ply count, lay-up sequence, and geometry—making it difficult to model deterministically. Thus, thermographic inspection in AFP is best framed as a passive thermography problem, where surface temperature irregularities are captured and analyzed to reveal potential anomalies.

Figure 3.1 illustrates the characteristic thermal footprints of key AFP defects us-

ing three-dimensional surface renderings. Freshly placed tows typically appear cooler than their surroundings due to transient thermal imbalance. From a thermographic perspective, gaps—where adjacent tows fail to contact—expose the underlying heated surface, producing localized hotspots. Missing tows exhibit a similar but broader thermal signature, reflecting the larger uncovered area. In contrast, overlaps create thicker laminate regions that attenuate heat flow, resulting in colder zones. Foreign bodies embedded within the lay-up obstruct thermal conduction and leave distinct localized depressions in the thermal profile. Tow splices—naturally occurring at tow junctions—introduce minor thickness discontinuities and produce thermal anomalies resembling those of foreign inclusions.

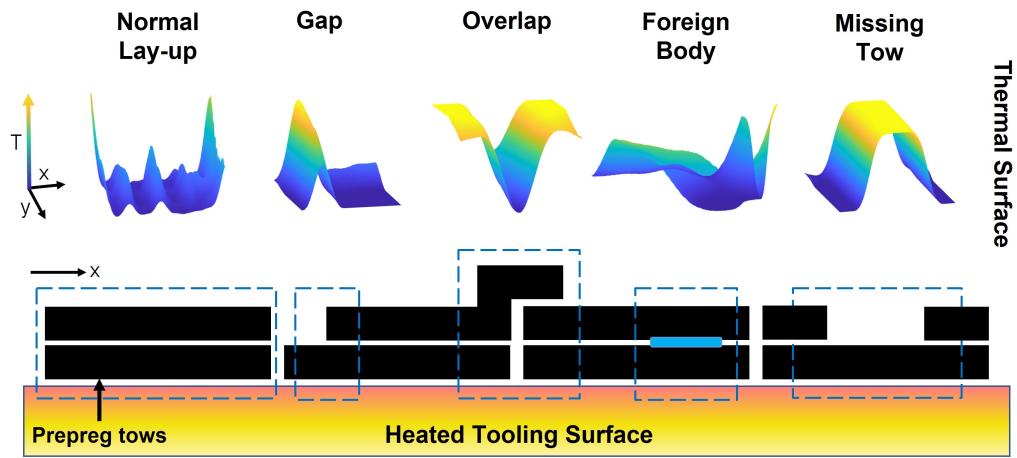


Figure 3.1 Thermal contrast patterns associated with common AFP defects.

### 3.2 In-Situ Thermal Vision Setup Design and Integration

A fully in-situ thermal vision system is developed and mounted directly on the AFP machine head to enable real-time thermographic inspection during lay-up. Unlike post-process or externally fixed setups, it provides continuous, unobstructed access to the heated region. The system comprises a thermally compliant sensor, a custom mechanical holder, and an integrated mounting scheme. Each component is detailed in the following subsections.

### 3.2.1 Sensor Requirements

Effective thermal inspection in AFP requires sensors capable of capturing high-resolution imagery at sufficient frame rates while remaining compatible with the process's thermal boundaries. In our system, spatial resolution governs the fidelity of qualitative defect cues for classification and segmentation, whereas frame rate controls the inter-frame overlap percentage used in motion estimation and lay-up reconstruction. To fulfil both spatial and temporal demands, we employ the FLIR A655sc uncooled thermal camera operating in the long-wave infrared (LWIR) band. It delivers a resolution of  $640 \times 480$  pixels and records at 50 fps in full-frame mode. While our system uses this native rate, the camera also supports windowed configurations up to 200 Hz, offering flexibility for high-speed scenarios if required.

The camera accommodates two standard operating ranges ( $-40^{\circ}\text{C}$ – $150^{\circ}\text{C}$  and  $100^{\circ}\text{C}$ – $650^{\circ}\text{C}$ ), with a measurement accuracy of  $\pm 2^{\circ}\text{C}$ —well suited to the typical lay-up surface range of  $20$ – $28^{\circ}\text{C}$ . A  $45^{\circ}$  wide-angle infrared lens ensures full coverage of the region of interest (RoI) while preserving geometric consistency. The camera is mechanically integrated using a custom-designed holder, described next.

### 3.2.2 Flexible Holder Mechanism Design

To overcome the limitations of post-process inspection, our system adopts an in-situ monitoring strategy requiring continuous and stable thermal imaging during AFP operation. This imposes strict demands on the positioning and orientation of the thermal camera relative to the region of interest (RoI). The design challenge lies in combining flexibility for visual alignment with structural rigidity under motion. To address this, the holder mechanism incorporates three mechanical degrees of freedom: adjustable height, adjustable arm length, and tunable viewing angle—each allowing precise camera alignment across varying part geometries.

The mechanism attaches to the right side of the AFP machine head via a circular squeeze clamp that secures a 14 mm cylindrical connection beam. This beam enables smooth vertical repositioning of the entire assembly. A base platform connected to the beam supports the thermal camera and houses a slider mechanism that adjusts arm length through a machined channel. Once the required extension is reached, a clamping plate locks the slider in place. Viewing angle control is achieved by interchanging the slider with pre-angled variants. These three adjustable axes form

the core of the mechanical contribution. A streamlined protective casing is added to reduce impact vulnerability and tidy cable routing. The complete design is shown in Figure 3.2.

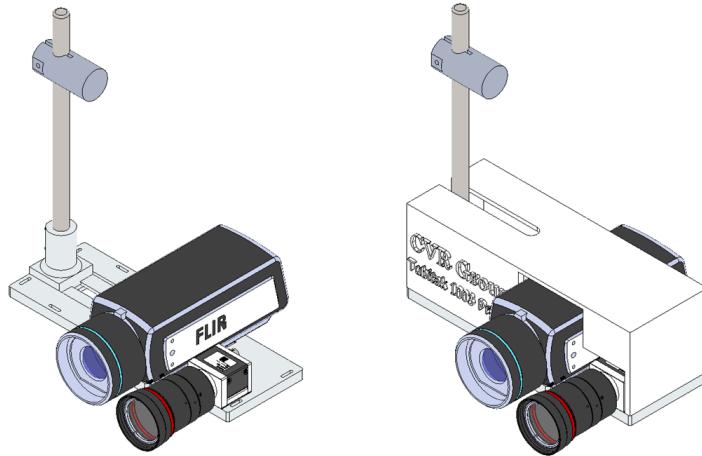


Figure 3.2 Assembled thermal vision holder mechanism: (Left) open functional configuration, (Right) final with protective casing.

### 3.2.3 Integration to AFP Machine Head

The mechanical components of the thermal vision system were fabricated using a mix of manual and digital manufacturing methods. The cylindrical connection beam, which endures the highest mechanical stress, was turned from hardened steel using a manual lathe. All remaining structural parts—including the base, slider, and clamp adapter—were CNC-milled from lightweight aluminium alloy to minimize load on the AFP head while ensuring rigidity. The protective casing was 3D-printed from white PLA to provide impact shielding and cable management.

Figure 3.3 shows the manufactured assembly without its casing. Final integration onto the AFP machine head is presented in Figure 3.4, where parameters were manually tuned through dry-run calibration. The final operational configuration used a height of 183.5 mm, arm length of 81.3 mm, and viewing angle of 10°.

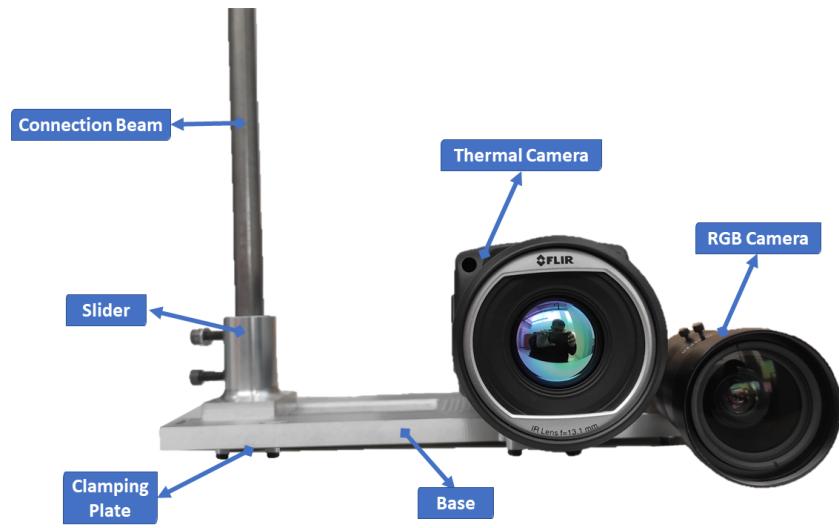


Figure 3.3 Manufactured thermal vision system without protective casing.

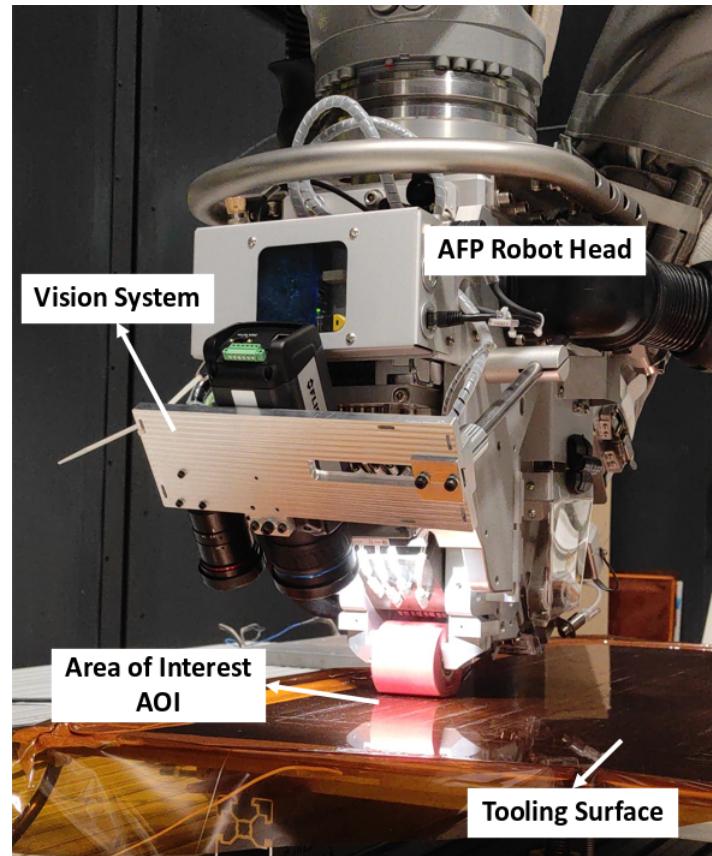


Figure 3.4 Integrated thermal vision system mounted on the AFP machine head.

### 3.3 Closed-Loop Thermal Conditioning System

Maintaining consistent surface temperatures across the tooling area is essential not only for bonding integrity during AFP lay-up but also for reliable thermal imaging. To this end, a closed-loop heating system was implemented, comprising two identical heating plates, a control unit, and eight thermocouples positioned around the lay-up region. The system stabilized the surface at a target temperature of 40°C, with fluctuations confined to the 35°C–45°C range—providing sufficient thermal contrast while remaining within the safe thermal limits of the prepreg material.

The effectiveness of thermal defect detection depends strongly on the quality of thermal contrast in the captured imagery. However, enhancing this contrast must be achieved without compromising process integrity. Prepreg materials are typically deposited at surface temperatures between 20°C and 28°C, and exceeding safe thresholds can degrade bonding quality and distort the fiber architecture. In addition, the tooling surface exhibits thermal non-uniformities, including geometric warping and edge–center temperature gradients of up to 7°C, as illustrated in Figure 3.5. While such variability complicates process control, it introduces visual diversity into the dataset. For instance, this diversity is strategically leveraged in the hierarchical defect identification module of the first proposed framework to improve generalization across defect types and process variants, whereas in the ThermoRAFT-AFP motion estimator, edge-induced variability is selectively removed through spatial cropping to optimize estimation accuracy and computational efficiency.

### 3.4 Thermal Camera Spatial Calibration

Camera calibration is essential for extracting metric information from thermal imagery, correcting lens-induced distortions, and ensuring geometric consistency for downstream tasks such as motion estimation, defect localization, and lay-up reconstruction.

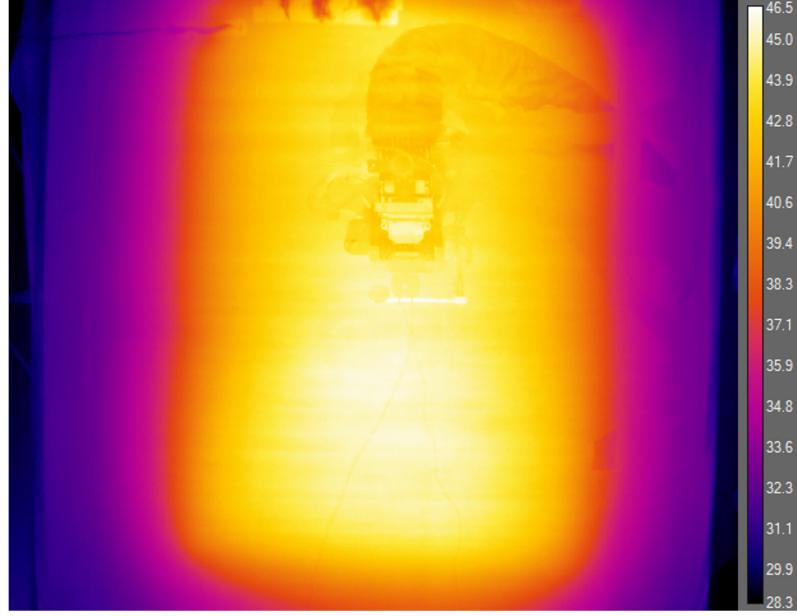


Figure 3.5 Tooling surface temperature map showing a  $+7^{\circ}\text{C}$  edge–center gradient caused by imperfect contact with heating plates.

### 3.4.1 Correspondence Problem

Camera calibration involves estimating intrinsic and extrinsic parameters by minimising the projection error between known 3D reference points and their corresponding 2D image locations. Given a set of correspondences  $\{P_W, p\}$  between world points  $P_W$  in a global coordinate frame ( $W$ ) and their projections  $p$  onto the image plane, the calibration task solves for the projection model:

$$(3.1) \quad \lambda p_i = M P_{W_i}, \quad M = K [R \ T]$$

Here,  $M$  is the full camera projection matrix,  $K$  is the  $3 \times 3$  intrinsic matrix, and  $[R|T]$  is the extrinsic transformation incorporating rotation and translation from world to camera coordinates. The intrinsic matrix  $K$  includes the focal lengths  $f_x, f_y$ , the optical center  $(u_0, v_0)$ , and the skew parameter  $s$ . A point  $P_W$  is thus projected to image coordinates  $p_i = [u_i, v_i, 1]^T$  via:

$$(3.2) \quad p_i = \begin{bmatrix} \frac{m_1 P_W}{m_3 P_W} \\ \frac{m_2 P_W}{m_3 P_W} \\ 1 \end{bmatrix}$$

Various methods exist for estimating  $K$ ,  $R$ , and  $T$ , such as the Direct Linear Transform (DLT) (Hartley & Zisserman (2003)), Heikkila’s approach (Heikkila & Silvén (1997)), and Zhang’s algorithm (Zhang (2000)). The latter is widely adopted for its robustness and ability to jointly estimate radial and tangential distortions. These distortion parameters are critical in the AFP input preprocessing pipeline, where correcting geometric distortions improves both image quality and downstream defect detection accuracy.

### 3.4.2 Calibration Setup and Results

Thermal calibration of infrared cameras requires tailored procedures due to the fundamental differences in sensing modality compared to visible-light cameras. Standard printed checkerboards, which offer high contrast for RGB calibration, exhibit poor radiative contrast in thermal images and thus fail to yield sharply defined corners under ambient conditions.

To overcome this, a mask-based calibration setup was designed, inspired by prior work (Hilsenstein (2005); Ng, Du & others (2005)). The custom calibration mask was cut from a heat-resistant composite sheet ( $320 \times 320 \times 3$  mm) using a KUKA KR 16-2C-F waterjet robot. A  $7 \times 11$  checkerboard pattern was machined into the plate, as shown in Figure 3.6. This mask was then mounted on a heated calibration rig comprising two heating plates, four metallic risers, and insulation layers, creating strong thermal gradients between mask holes and surrounding material.

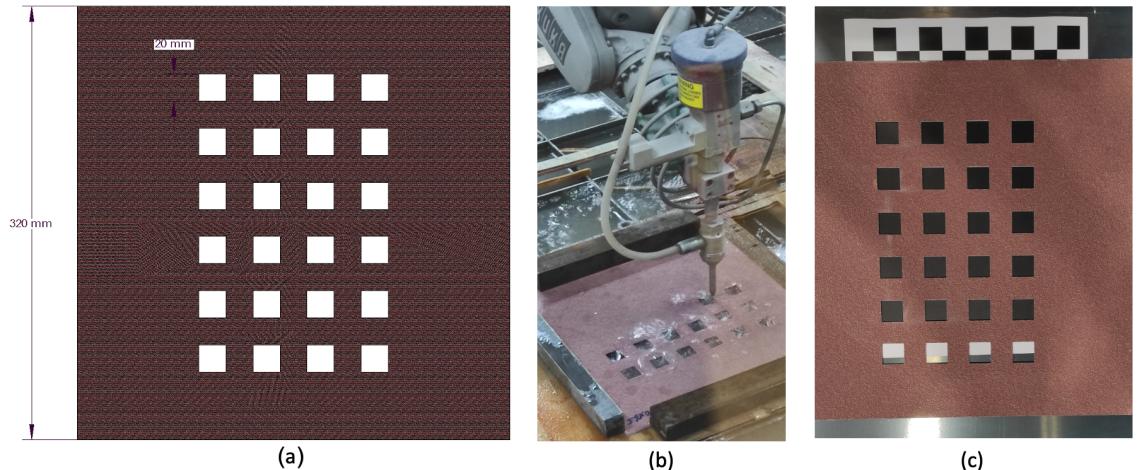


Figure 3.6 Thermal calibration mask: (a) CAD design, (b) waterjet cutting process, (c) manufactured composite mask.

The mask was integrated into a thermal calibration rig comprising four elevation

mounts, a flat aluminium heating plate, and underlying thermal insulation. The setup uses the closed-loop heating system introduced earlier, enabling stable high-contrast thermal patterns by heating the background to 150°C while the air-cooled mask regions remained cooler. Figure 3.7 shows the fully assembled rig with component labels and its thermal response during operation.

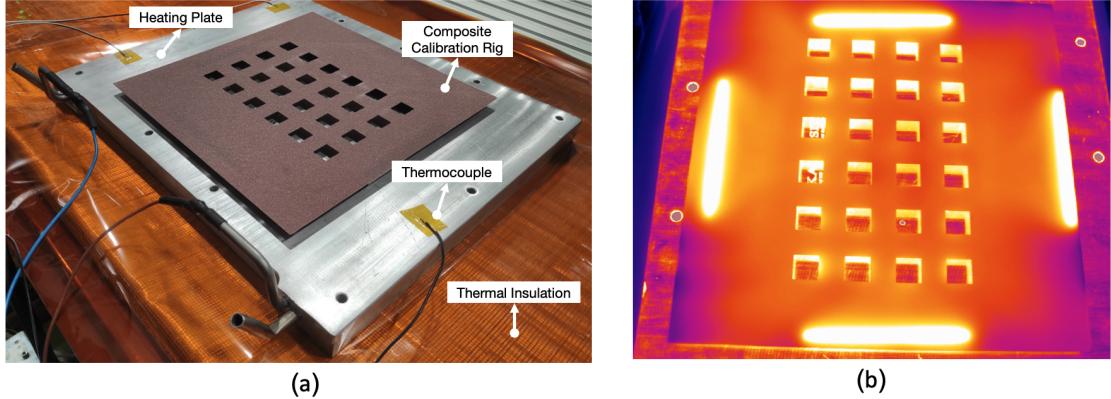


Figure 3.7 Thermal calibration setup: (a) Assembled rig with labeled components; (b) thermal image showing strong mask–background contrast at 150°C.

Twelve thermal images were captured from various viewpoints and processed using Zhang’s algorithm. Detected corner points were matched with their 3D mask coordinates and used to solve for the intrinsic matrix  $K$ , radial distortion  $d_{\text{rad}}$ , and tangential distortion  $d_{\text{tan}}$  with subpixel precision:

$$(3.3) \quad K = \begin{bmatrix} 802.37 & 0 & 327.29 \\ 0 & 801.08 & 237.58 \\ 0 & 0 & 1 \end{bmatrix}$$

$$(3.4) \quad d_{\text{rad}} = \begin{bmatrix} -0.2219 & 0.2881 & -0.001 \end{bmatrix}$$

$$(3.5) \quad d_{\text{tan}} = \begin{bmatrix} -0.0005 & 0 \end{bmatrix}$$

The 3D poses of the calibration images relative to both the camera and mask coordinate systems are shown in Figure 3.8. These extrinsic parameters, encoding the camera’s position and orientation per view, enable metric retrieval from image-plane to world coordinates and are subsequently reused in lay-up reconstruction and global visualisation workflows.

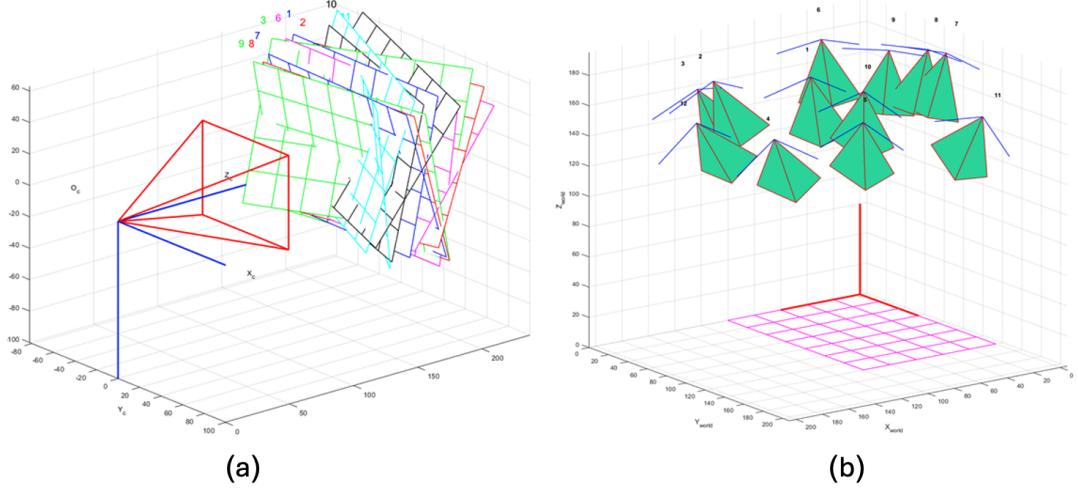


Figure 3.8 Thermal calibration extrinsics: (a) camera-centered view; (b) world-centered view.

### 3.5 Experimental Setup and Data Acquisition Protocol

All experiments were conducted on a *Coriolis C1* AFP machine equipped with a robotic head capable of placing up to eight 6.35 mm-wide tows. The processing pipeline was executed on a dedicated workstation featuring an Intel Xeon W-2275 CPU (14-core @ 3.30 GHz, Turbo up to 4.60 GHz), 32 GB DDR4 RAM, and an NVIDIA Quadro RTX 5000 GPU. All tasks were implemented in MATLAB and Python using widely adopted libraries (Bradski (2000); Paszke, Gross, Chintala, Chanan, Yang, DeVito, Lin, Desmaison, Antiga & Lerer (2017); Van der Walt, Schönberger, Nunez-Iglesias, Boulogne, Warner, Yager, Gouillart & Yu (2014)).

A broad spectrum of process and environmental conditions was covered, as summarised in Table 3.1. The parameter set includes varying lay-up speeds, ply lengths, fiber orientations, layer counts, and defect scenarios—reflecting realistic production variability and enabling robust performance testing across both defect inspection and reconstruction tasks.

Thermal data acquisition was performed asynchronously using the camera’s full-frame mode at 50 Hz, ensuring maximum spatial fidelity for high-precision defect analysis and segmentation. This configuration produced high inter-frame overlap, with pitch sizes ranging from 2 mm to 6 mm across the tested lay-up speeds—crucial for accurate motion-based reconstruction. By operating independently of the AFP controller, the acquisition protocol supports a major objective of this work: enabling a fully automated, machine-independent inspection pipeline driven solely by the thermal camera and the proposed frameworks. However, excessive redundancy can

Table 3.1 AFP Process and Environmental Conditions

Parameter	Value / Range
Tow width	6.35 mm
Tow thickness	160 $\mu$ m
Material	Thermoset
# Tows / Layers	4 / 1–4
Orientations	[0°, 45°, 90°]
Lay-up speed	0.1–0.3 m/s
Course length	300–800 mm
Compaction pressure	3–6 bar
Tool temp / Material temp	32–50 °C / 20–28 °C
Ambient temperature	19–24 °C

degrade the efficiency of detection and classification models.

To address this, the acquired stream was repurposed into two datasets tailored to the operational needs of each proposed framework: (1) a large-scale, full-rate dataset used in the *ThermoRAFT-AFP* motion estimation and lay-up reconstruction framework, and (2) a subsampled, minimal-overlap subset used for training and validating defect detection, classification, and segmentation models. This dual-dataset strategy maximises data utility and enables consistent benchmarking across both frameworks. Technical specifications of these datasets are detailed in Chapter 6.

### 3.6 Synthetic Defect Scenarios

The hybrid frame-wise defect analysis and quality assessment framework developed in this work depends on robust learning-based models capable of detecting, classifying, and segmenting a wide spectrum of AFP defects. Effective training and generalization require datasets that cover not only common anomalies but also rare, challenging cases across diverse conditions. However, many defect types—particularly those from extreme deviations or infrequent events—are underrepresented in naturally acquired data. To address this, we designed a structured set of synthetic defect scenarios to augment the dataset while preserving realism and full in-situ compatibility.

Three main categories of synthetic scenarios were implemented: (1) missing tows, (2) artificial gaps and overlaps, and (3) foreign bodies. Missing tow conditions, including

isolated and multi-tow absences, were introduced by selectively omitting tows in the lay-up design across up to four layers. These defects were embedded during normal operation and represent a structurally critical class. Generating artificial gaps and overlaps for in-situ thermal monitoring posed a greater challenge, as post-process methods are unsuitable for dynamic AFP workflows. To simulate them, prepreg tows were modified on the spool using cutting or folding techniques. The resulting discontinuities were introduced naturally during lay-up. Overlaps arose through fold-induced buckling or manual insertion between layers.

The foreign body category was extended to include realistic contamination scenarios and FOD—small consumables accidentally embedded into the laminate. Materials included poly backing film, Kapton tape, vacuum bag scraps, carbon fiber debris, and heat-resistant mesh. These were shaped irregularly and randomly introduced onto the tow or surface before compaction to mimic unintentional inclusions. Tow splice defects—naturally occurring at junctions between prepreg roll segments—were also captured during experiments. These regions cause localised thickness variation and are thermally distinct due to altered conductivity. While some high-end AFP systems detect them via encoded markers, this feature is not widespread. Including splices thus improves industrial relevance and operator-assistive capabilities.

All synthetic scenarios were integrated into the dataset pipeline and annotated for supervised training, validation, and performance benchmarking of detection, classification, and segmentation models. Their inclusion boosts both diversity and robustness of the learning-centric inspection system.

### 3.7 Universal Preprocessing Pipeline

The performance of the proposed inspection frameworks is highly dependent on consistent and reliable image preprocessing. Raw thermal images may exhibit radiometric noise, geometric distortion, and non-uniform viewpoint, all of which reduce the accuracy of learning-based and geometric analysis modules. To mitigate these issues, we introduce a unified preprocessing pipeline applied to every captured frame prior to use in any downstream module. This standardized sequence ensures spatial and temporal alignment across all inputs, enabling clean, interpretable results throughout the system.

Each thermal image undergoes four key steps, as illustrated in Figure 3.9. First, 16-

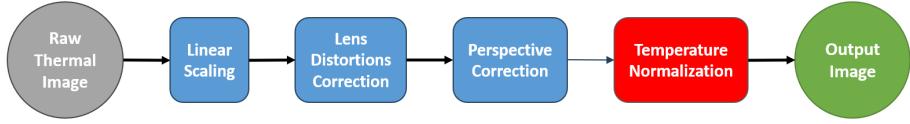


Figure 3.9 Flowchart of the universal preprocessing pipeline.

bit radiometric measurements are linearly scaled to 8-bit grayscale values to reduce computational cost while retaining critical thermal contrast. Second, lens-induced geometric distortions are corrected using intrinsic parameters estimated during the thermal camera calibration phase.

The third step performs spatial alignment via homography-based perspective correction. A quadrilateral ROI is first defined to isolate the region of maximum thermal contrast—typically just beneath the compaction roller—and remove background clutter. Among the two main preprocessing strategies—processing skewed frames or applying geometric rectification—we adopt the latter, as it produces a spatially standardized input domain better suited for geometric analysis and learning-based inference. Moreover, it simplifies many subtasks in the system such as tow identification and motion estimation problems by aligning the lay-up direction vertically, reducing the tasks to a simpler 1D version. A homography transformation then restores a fronto-parallel view aligned with the tool surface, preserving geometric properties such as collinearity and scale for accurate downstream processing. The applied transformation is given by:

$$(3.6) \quad \lambda \begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

where  $h_{ij}$  are the elements of the  $3 \times 3$  homography matrix that transforms a point  $(x, y)$  into  $(x', y')$ , and  $\lambda$  is a scaling factor. The final stage addresses temporal thermal inconsistencies across the lay-up course. Older regions exposed to compaction heat appear warmer than freshly deposited areas, producing a vertical gradient. To correct this, the average temperature across each image row is computed and subtracted, yielding a temporally normalised map with enhanced defect visibility. As shown in Figure 3.10, features such as overlaps and foreign objects become more distinct after this operation.

To illustrate the combined effect of the full pipeline, Figure 3.11 compares a raw thermal image of a flawless lay-up with its preprocessed counterpart. The resulting output is both spatially and radiometrically standardized, ready for use across all modules. Unlike task-specific preprocessing routines, this universal design ensures interoperability across all modules, simplifying deployment and system integration.

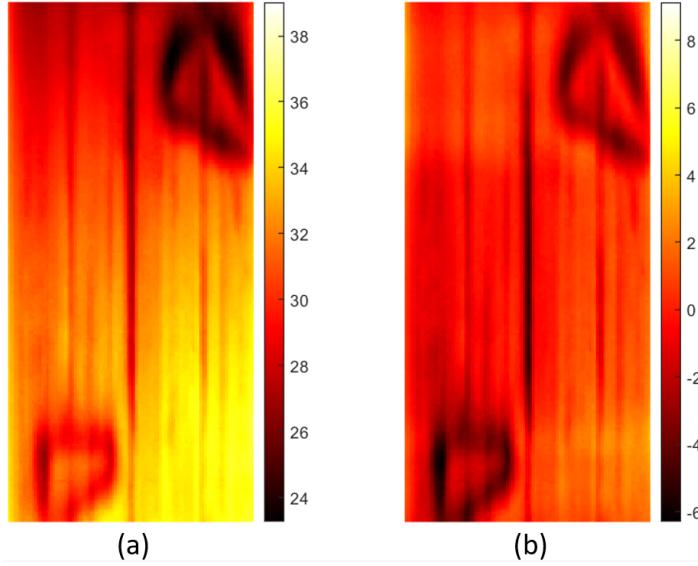


Figure 3.10 Temperature normalisation: (a) perspective corrected image; (b) temperature-normalised version showing deviation from row-wise mean.

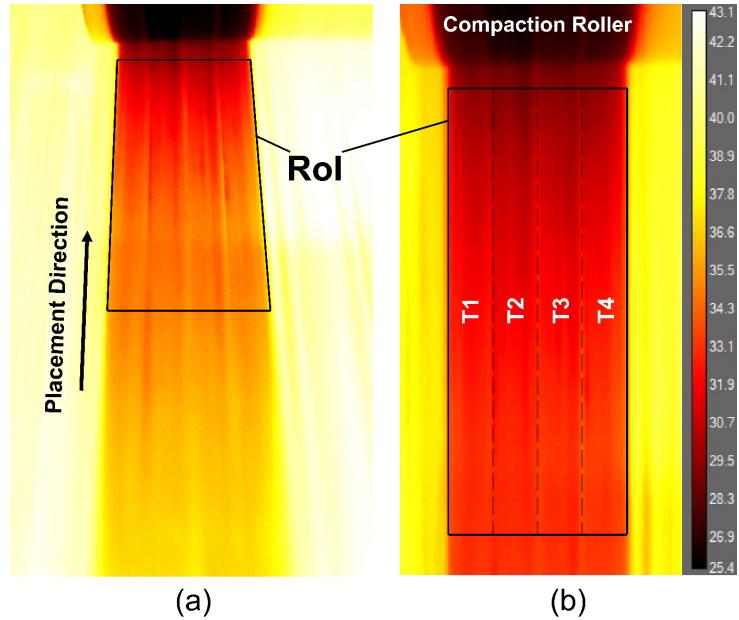


Figure 3.11 (a) Raw thermal image of a flawless lay-up; (b) corresponding dataset-grade preprocessed result.

Figure 3.12 presents representative thermal images of common AFP defect types studied in this work, shown after the full preprocessing pipeline. These include gaps (G), missing tows (MT), overlaps (O), tow splices (TS), and FOD (FB)—each showing distinct thermal patterns that support reliable interpretation. Spatial standardization and temperature normalisation enhance the saliency and separability of these defects, enabling accurate detection, classification, and segmentation within the proposed learning-centric inspection framework. These standardized images are consumed by both frameworks at different rates: sparsely for frame-wise defect

analysis (Framework 1) and densely for thermal motion tracking (Framework 2).

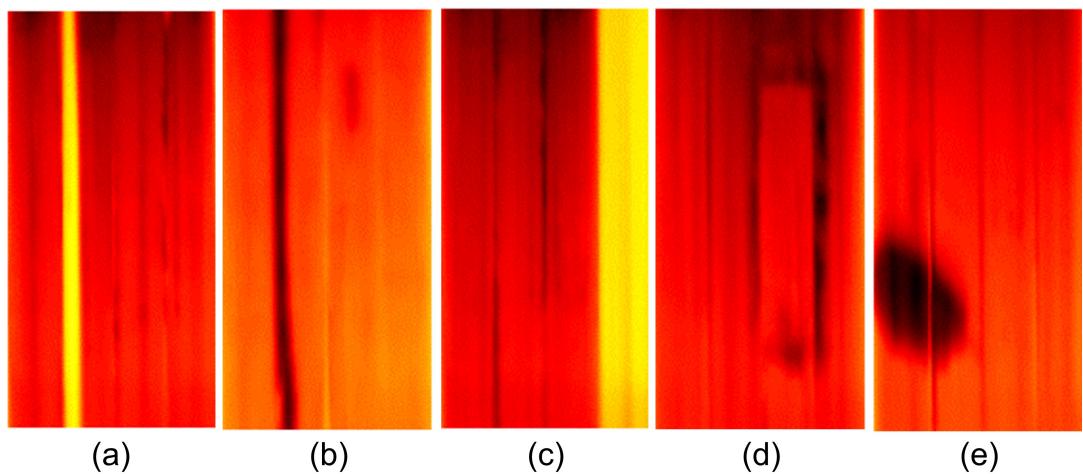


Figure 3.12 Examples of preprocessed defective lay-up instances: (a) gap (b) overlap (c) missing tow (d) tow splice and (e) foreign body.

## 4. HYBRID FRAME-WISE DEFECT ANALYSIS AND QUALITY ASSESSMENT FRAMEWORK

This chapter presents the first core component of the proposed inspection system: a hybrid, frame-wise analysis framework that detects, classifies, and evaluates defects in thermal images acquired during the AFP process. This chapter incorporates material adapted from the author's peer-reviewed journal publication (Zemzemoglu, Unel & Tunc (2024)) reused with minor modifications, including the framework design, dynamic tow identification module, lay-up quality evaluation module, and associated figures. The framework integrates hybrid learning algorithms and traditional computer vision in a hierarchical cascade, enabling real-time operation, tow-level reasoning, and quality assessment using only a single thermal camera. Its outputs form the foundation for downstream decision-making and are later contextualized within a global, motion-aware framework introduced in the following chapter.

### 4.1 Framework Overview

The architecture of the proposed framework addresses key challenges in in-situ AFP process monitoring and quality inspection by combining modularity, real-time operability, and hybrid algorithmic design. The framework adopts a multi-level paradigm that tailors task-specific solutions to the nature, complexity, and data constraints of each processing stage. This design facilitates a judicious integration of model-based and learning-based methods—reserving intensive learning algorithms for scenarios where they are most effective, while relying on traditional vision techniques where data scarcity or simplicity permits.

As shown in Figure 4.1, the system comprises three tightly coupled modules: **Dynamic Tow Identification**, **Hierarchical Defect Identification**, and **Lay-up Quality Evaluation**. These modules operate in parallel or conditionally, contribut-

ing to system responsiveness and computational efficiency. To ensure robustness and adaptability, the design emphasizes generalizability across different AFP setups and manufacturing contexts, and prioritises interpretability to support human-in-the-loop decision making.

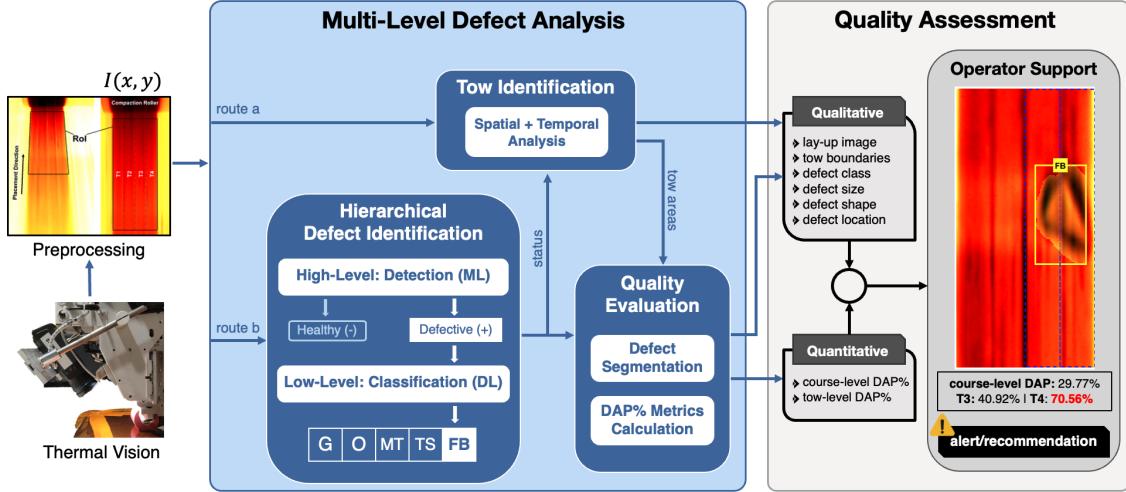


Figure 4.1 Architecture of the proposed Hybrid Frame-wise Defect Analysis and Quality Assessment Framework, illustrated using a sample FOD defective thermal input.

The quality of AFP manufacturing is tightly linked to the precise placement of individual tows. The first module, **Dynamic Tow Identification**, performs setup-independent estimation of tow boundary positions using a custom spatial-temporal analysis algorithm. It outputs individual tow regions, which are subsequently referenced by other modules to facilitate tow-level reasoning and impact analysis.

The second module, **Hierarchical Defect Identification**, is central to the system's primary decision pipeline. It implements a two-stage classification strategy to assess the health of each thermal frame. A high-level binary classifier first determines whether the lay-up instance is healthy (-) or defective (+). If the instance is deemed defective, a low-level multi-class classifier is triggered to identify the specific defect class from a predefined set. This cascaded structure addresses the natural data imbalance—where healthy cases dominate—and reduces unnecessary computation by bypassing the multi-class stage for non-defective frames. Furthermore, this module operates in parallel with tow identification, allowing concurrent processing and contributing to real-time performance.

When a defect is detected, the final module—**Lay-up Quality Evaluation**—is conditionally activated. It applies a pixel-level segmentation algorithm to localise the defective region and extract key morphological features such as shape, size, and position. These features, combined with tow geometry from the first module, are

used to compute a novel metric called *DAP*, which quantifies the impact of each defect in relation to the effective lay-up area.

To meet industrial standards and achieve state-of-the-art performance, the system must fulfill:

- **Tow Identification Robustness:** Sub-pixel accuracy with average error  $< 1 \text{ px}$  and stable convergence within  $\pm 2 \text{ px}$ , even in the presence of defects and without prior assumptions.
- **Detection Reliability:** SVM-based classifier achieving  $\geq 95\%$  test accuracy and  $> 95\%$  recall, minimizing the likelihood of undetected defective lay-up instances.
- **Classification Accuracy:** DCNN achieving  $\geq 95\%$  overall and per-class accuracy across all studied defect types.
- **Segmentation Precision:** Mean pixel accuracy (PA)  $\geq 90\%$  and Intersection-over-Union (IoU) score  $\geq 0.65$ —yielding spatially conservative yet robust defect localization.
- **Real-Time Operation:** Minimum 5 fps throughput on standard CPU hardware across all modules: tow identification, detection, classification, and segmentation.

The synergistic integration of these modules forms the core of an AFP-DSS. Operating solely with a thermal camera and a lightweight algorithmic stack, the framework delivers real-time, frame-wise defect analysis enriched with both qualitative and quantitative insights. It enables precise identification, characterisation, and impact assessment—empowering operators with timely, data-driven guidance. When paired with the Motion-Aware Global Inspection Framework, these outputs gain temporal traceability, machine-independent alignment, and enhanced visual coherence across the evolving lay-up surface—paving the way for a fully automated, end-to-end, and operator-assistive thermographic inspection solution.

## 4.2 Dynamic Tow Identification

This module enables frame-wise identification of individual tows in thermal images by automatically and reliably detecting the boundaries between adjacent tows. In fixed-stiffness AFP manufacturing, tow boundaries correspond to narrow vertical gaps that separate neighbouring tows. These appear as low-temperature lines in thermographic imagery due to slight heat dissipation across inter-tow regions. Leveraging this phenomenon, the module fuses thermal cues with targeted computer vision techniques to infer tow geometry.

Figure 4.2 illustrates the thermal basis underpinning this approach. Panel (a) shows a 3D surface plot of a lay-up course, with pixel intensity mapped to temperature. After Gaussian smoothing ( $\sigma = 0.5$ ), elongated thermal valleys emerge that align visually with inter-tow gaps. Panel (b) presents equidistant cross-sectional profiles in which local minima consistently correspond to tow boundaries, highlighted with transparent vertical planes. This empirical analysis supports the use of thermal gradients as reliable indicators for boundary localization.

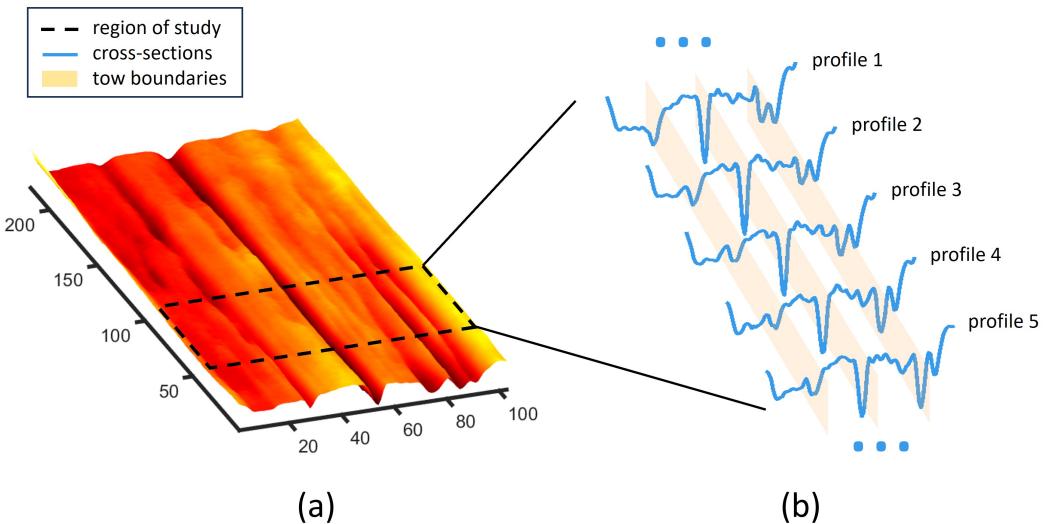


Figure 4.2 Tow boundary analysis in thermographic AFP inspection: (a) 3D thermal surface map with inter-tow valleys; (b) cross-sectional intensity profiles showing local minima aligned with tow boundaries.

Accurate tow boundary detection remains a nontrivial challenge due to thermal

variability, geometric distortion, and shifting setup conditions. While some methods process raw, skewed frames and others employ geometric rectification to simplify the spatial structure, both rely on assumptions—such as fixed tow layouts or static camera configurations—that may not hold in practice. These limitations motivate the need for a robust, setup-independent solution that performs reliably under real-time, frame-wise constraints.

The proposed solution integrates classical computer vision tools into a lightweight and modular pipeline. Preprocessing first produces rectified top-down views, reducing the two-dimensional search space to a one-dimensional estimation task along the x-axis. As a result, the algorithm estimates a series of x-coordinates corresponding to inter-tow separations. A spatial-temporal analysis algorithm then tracks tow boundary positions over time. Exploiting the sequential nature of thermal data, the algorithm incrementally refines boundary predictions across frames with minimal computational overhead. Estimates rapidly converge to stable x-coordinates, enabling accurate, consistent, and efficient tow extraction throughout the lay-up sequence. The algorithmic formulation is presented in the following subsection.

#### 4.2.1 Spatial-Temporal Analysis Algorithm

The algorithm operates in two stages: a spatial estimation followed by temporal refinement. First, each incoming frame is analysed in isolation to detect potential tow boundaries based on image features. Then, these estimates are refined using knowledge accumulated from previous frames, improving stability and suppressing outliers. This structure leverages the sequential nature of AFP imagery, offering resilience against noise and frame-level inconsistencies. The full procedure is summarised in Algorithm 1.

To enhance edge detection and reduce noise, optional filtering is applied prior to edge extraction. Three settings are supported: no filter, a linear filter, or a nonlinear filter. The linear option uses a vertical Scharr operator, an enhanced gradient filter offering strong edge responses and rotational symmetry (Jähne, Scharr, Körkel, Jähne, Haußecker & Geißler (1999)). The nonlinear option applies a morphological Top-hat filter, which enhances vertical line features by subtracting local background using erosion–dilation cycles (Gonzalez (2009)). While not originally designed for line detection, its sensitivity to local contrast suits thermographic imagery. The *filterType* parameter allows flexible adaptation to varying imaging conditions.

---

**Algorithm 1** Tow identification via online spatial-temporal analysis

---

```

Initialise prevXests to store boundary history
for each frame in imgSeq do
    Apply filterType filtering
    edges  $\leftarrow$  detect edges using edgeMethod
    vHist  $\leftarrow$  sum edges along vertical axis
    for k from mDist to length(vHist)  $- mDist do
        if vHist[k] is local max and  $> T$  then
            peaks  $\leftarrow$  k
        currXest  $\leftarrow$  merge peaks closer than mDist
        if not first frame then
            updatedXest  $\leftarrow$  average currXest with prevXests
        else
            updatedXest  $\leftarrow$  currXest
        Append updatedXest to prevXests
        Compute estimation error$ 
```

---

Edge maps are extracted using the Canny detector, chosen for its precision and continuity-preserving characteristics (Canny (1986)). Non-maximum suppression and hysteresis thresholding isolate high-confidence contours. These edges are projected along the vertical axis to form a one-dimensional histogram *vHist*, transforming the problem into a line-wise peak detection task. Local maxima exceeding threshold *T* within a  $\pm mDist$  window are selected as candidate peaks. Closely spaced peaks are then merged to eliminate boundary duplication caused by image artifacts.

Temporal refinement is achieved by smoothing each current estimate using previous frame history. We evaluate three averaging strategies: simple moving average ( $SM_{avg}$ ), weighted moving average ( $WM_{avg}$ ), and cumulative average ( $C_{avg}$ ). Each method produces a refined x-coordinate  $x_t$  for a given tow boundary as follows:

$$(4.1) \quad SM_{avg}(t) = \frac{1}{k} \sum_{i=t-k+1}^t x_i$$

$$(4.2) \quad WM_{avg}(t) = \frac{\sum_{i=t-k+1}^t w_i x_i}{\sum_{i=t-k+1}^t w_i}$$

$$(4.3) \quad C_{avg}(t) = \frac{1}{t} \sum_{i=1}^t x_i$$

where *k* is the averaging window size and  $w_i$  are assigned weights. At each timestep *t*, updated boundary positions are recorded, and estimation errors are computed

as the Euclidean distance to manually annotated ground truth. The resulting x-coordinates are passed to downstream modules for further use in segmentation and quality assessment tasks.

### 4.3 Hierarchical Defect Identification

The second core module of the proposed framework is dedicated to identifying both the presence and type of defects in each thermographic frame acquired during the AFP process. As discussed in the system overview, this task is implemented using a two-stage hierarchical classification pipeline. The first stage determines whether a given frame contains a defect or not. If a defect is detected, a second stage is triggered to classify its type. This hierarchical arrangement allows efficient allocation of computational resources by avoiding unnecessary processing on healthy frames, while also addressing the natural class imbalance often encountered in AFP datasets, where defective cases are relatively rare. The hierarchical design, deep learning architecture, and associated figures in this section are adapted from the author’s peer-reviewed conference publication (Zemzemoglu & Unel (2023)), in accordance with corresponding reuse policies.

#### 4.3.1 High-Level Defect Detector

The first level of inspection focuses on binary classification: determining whether a given frame is ‘Healthy’ (H) or ‘Defective’ (D). Existing strategies for this task can be broadly categorized into model-based and data-driven approaches. Model-based methods—such as thresholding, rule-based masking, or manually tuned filters—can offer reliable performance in static, well-controlled environments. However, their generalizability to dynamic AFP conditions remains limited, especially when confronted with variations in thermal background, camera angle, or defect morphology.

In contrast, data-driven techniques, particularly MLs, demonstrate greater robustness and adaptability to real-world scenarios. Among these, classical MLs models are especially suitable for highly separable problems such as this one, where thermal differences between healthy and defective lay-ups are generally pronounced. Deep

learning models, while powerful, are often unnecessary at this level due to their high data and computational demands. Classical models strike an ideal balance: they require modest training data, operate with low latency, and perform reliably when supported by well-crafted feature representations.

#### 4.3.1.1 Gabor-Based Feature Extraction

The classifier’s performance depends heavily on the quality of the input features. For thermal imagery in AFP, where many defects manifest as local texture anomalies, texture descriptors are particularly effective. Several studies have explored diverse feature types—including geometric, motion-based, and color-derived attributes—but texture features remain the most discriminative in fibrous composite surfaces (Mahajan, Kolhe & Patil (2009)). Among texture-based methods, the Gabor Transform has proven especially effective for capturing directional and frequency-specific information in fabric-like materials (Hanbay, Talu & Özgüven (2016)).

Turner (Turner (1986)) was among the first to apply GFs to texture analysis, and their advantages have since been validated in numerous comparative studies, particularly for fabric defect inspection (Javed, Mirza & others (2013)). Consequently, a GF-based pipeline is adopted for feature extraction in our high-level detector. These features are also reused by the lay-up quality evaluation module, significantly reducing redundant computation.

Once extracted, these GFs-based responses form a compact, low-dimensional representation that is passed to the classifier. Mathematically, a 2D Gabor filter is defined as:

$$(4.4) \quad g(x, y; \lambda, \theta, \psi, \sigma, \gamma) = \exp\left(-\frac{x'^2 + \gamma^2 y'^2}{2\sigma^2}\right) \cos\left(2\pi \frac{x'}{\lambda} + \psi\right)$$

$$(4.5) \quad x' = x \cos(\theta) + y \sin(\theta), \quad y' = -x \sin(\theta) + y \cos(\theta)$$

Here,  $\lambda$  represents the filter wavelength,  $\theta$  its orientation,  $\psi$  the phase offset,  $\sigma$  the scale of the Gaussian envelope, and  $\gamma$  the aspect ratio. A filter bank is constructed by sweeping  $\lambda$  and  $\theta$ , enabling extraction of multi-scale, multi-orientation texture responses from the input image. These responses are subsequently downsampled and

normalized to form a feature vector for classification, as visualised in Figure 4.3.

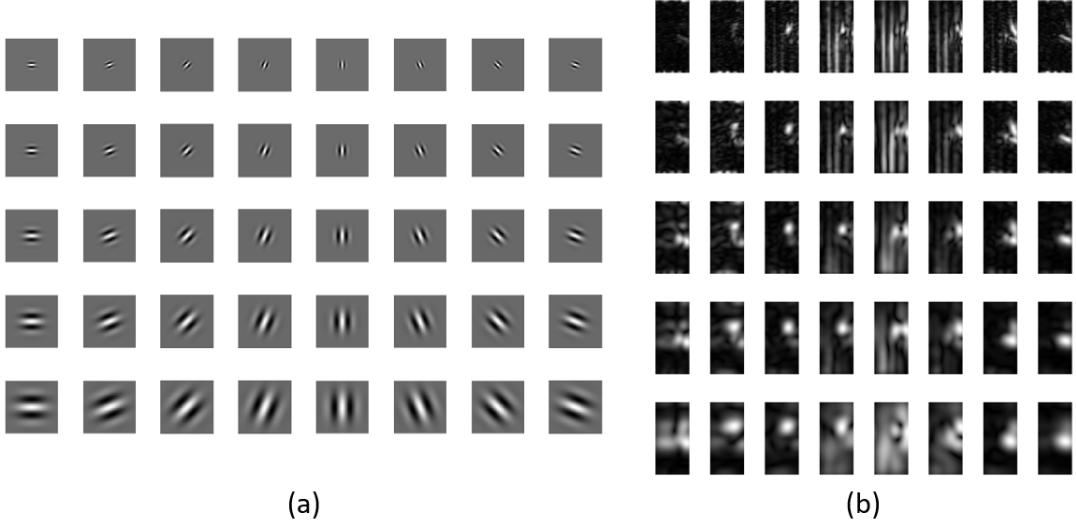


Figure 4.3 Gabor filter-based feature extraction: (a)  $5 \times 8$  filter bank; (b) sample convolution outputs from a defective frame.

#### 4.3.1.2 SVM Classifier Model

Once the feature vector is computed, it is passed into a lightweight classifier for final decision-making. Various machine learning algorithms were considered for this task. While Naive Bayes and Random Forests perform well on structured data, they are less effective for high-dimensional handcrafted features. In contrast, SVMs are particularly robust for binary classification involving dense, texture-rich descriptors such as those extracted using GFs. Their ability to maximize the margin between classes supports strong generalization performance—critical in dynamic industrial contexts (Cervantes, Garcia-Lamont, Rodríguez-Mazahua & Lopez (2020)). Figure (4.4) offers an intuitive view of how the SVM forms soft-margin boundaries in feature space.

SVMs also offer a key advantage in scenarios with limited annotated data, which is often the case in specialized manufacturing pipelines. These traits led to the adoption of a soft-margin SVM formulation in this study. This approach introduces a slack variable  $\xi_i$  for each training instance, allowing the model to tolerate bounded misclassifications while maintaining a maximized margin. A regularization term  $C$  manages the trade-off between generalization and empirical accuracy. The

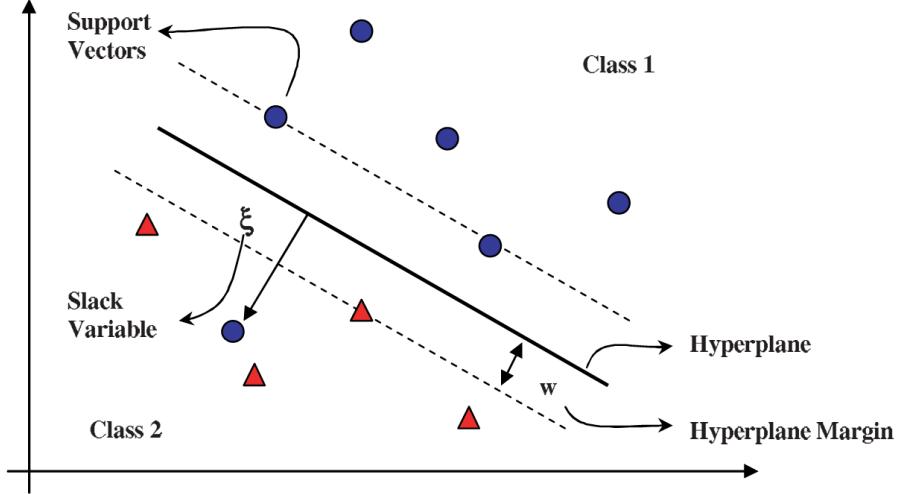


Figure 4.4 Soft-margin SVM decision boundary separating healthy and defective classes in feature space.

formulation is:

$$(4.6) \quad \min \frac{1}{2} \|\mathbf{w}\|^2 + C \sum_{i=1}^l \xi_i$$

$$(4.7) \quad \text{s.t. } y_i(\mathbf{w}' \mathbf{x}_i - b) \geq 1 - \xi_i \quad \forall i = 1, \dots, l$$

Here,  $\mathbf{w}$  is the weight vector,  $b$  is the model bias,  $C$  is the penalty term, and  $\xi_i$  are slack variables.

To effectively model nonlinear decision boundaries common in thermal imagery, the classifier uses the kernel trick. Kernel functions project the input space into a higher-dimensional feature space without explicitly computing the transformation, enhancing expressiveness while maintaining efficiency. We adopt the RBF kernel:

$$(4.8) \quad K(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\gamma \|\mathbf{x}_i - \mathbf{x}_j\|^2)$$

where  $\gamma$  adjusts the influence range of each support vector—lower values produce smoother boundaries. The model is trained on GF-derived features using grid search to optimize  $C$  and  $\gamma$ . Once deployed, it classifies incoming frames as either 'Healthy' (H) or 'Defective' (D) in real time. Frames flagged as defective are passed to the next inspection stage for detailed classification. As new defect instances arise during production, the training database can be expanded incrementally—supporting the system's long-term adaptability.

### 4.3.2 Low-level Defect Classifier

This second level of the inspection pipeline is invoked only when the high-level classifier identifies a frame as defective. It refines the classification by determining the specific defect type. The corresponding preprocessed thermal image is passed to a dedicated deep learning module that assigns it to one of five predefined defect categories. Since this stage requires richer feature understanding and greater representational power, a DCNNs is employed. The adopted network is designed to balance classification accuracy with runtime efficiency, ensuring suitability for near real-time deployment.

#### 4.3.2.1 Deep Learning Network Architecture

Designing an effective deep network requires careful calibration of architectural parameters—such as the number of layers, convolutional kernel sizes, pooling strategies, and activation functions. Several candidate architectures were examined, ranging from shallow to very deep configurations. However, given the trade-off between depth and computational load, a moderate-depth CNN was selected. This design captures the necessary visual complexity while maintaining efficient inference and manageable training demands.

The final architecture, shown in Table 4.1, consists of twelve layers: alternating convolution and max-pooling layers with gradually increasing kernel sizes, followed by two fully connected layers. All convolutional outputs are normalized and passed through ReLU activations, while the final dense layer employs the *Softmax* function to generate a probability distribution across the five defect classes.

Table 4.1 Optimized DCNN for low-level defect classification

Layer	Type	Kernel Size	# Filters
Input	Input Layer	[220 110 1]	—
1–2	Conv + MaxPool	[3×3], [2×2]	8
3–4	Conv + MaxPool	[5×5], [2×2]	8
5–6	Conv + MaxPool	[7×7], [2×2]	16
7–8	Conv + MaxPool	[9×9], [2×2]	32
9–10	Conv + MaxPool	[15×15], [2×2]	64
11	Dense Layer	300	—
12	Dense Layer (Softmax)	5	—

#### 4.3.2.2 Model Training and Refinement

Training the CNN involves optimizing the network weights to minimize the classification loss over the training dataset. To this end, the SGDM algorithm is adopted. SGDM accelerates convergence by smoothing the gradient descent path using a momentum term—a moving average of past gradients. The weight update rules are defined as:

$$(4.9) \quad V_t = \beta V_{t-1} + (1 - \beta) \nabla_W L(W, X, y)$$

$$(4.10) \quad W_t = W_{t-1} - \alpha V_t$$

where  $W$  represents the network weights,  $L$  is the loss function,  $\alpha$  is the learning rate,  $\beta$  is the momentum coefficient bounded by  $[0, 1]$  interval, and  $V_t$  is the velocity vector controlling update smoothness.

To improve training efficiency and reduce memory usage, mini-batch processing is used, with each batch consisting of 20 images. Hyperparameter tuning is performed empirically and reported in the Results chapter. After training, the model is validated using held-out data and tested on previously unseen samples to assess its generalization. Once deployed, this module enables fine-grained defect classification in real time, thereby enhancing the interpretability and traceability of the AFP inspection system.

## 4.4 Lay-up Quality Evaluation

The final stage of the proposed inspection framework focuses on a comprehensive quality assessment of thermal frames flagged as defective. Upon detection of a defect by the high-level classifier, this evaluation module is triggered to both localize the anomalous region and quantify its severity. The core of this module consists of two sequential components: a segmentation algorithm that delineates the spatial extent of the defect, and a metric-based evaluation that translates these regions into actionable quality scores. This dual approach enables the system to transition from binary defect presence to a nuanced frame-wise defect impact assessment—both visually and numerically—thus bridging detection with repair-oriented decision support.

#### 4.4.1 Defect Segmentation and Localization

This component addresses the problem of identifying and isolating defect regions within a defective input frame. Formulated as a binary segmentation task, the goal is to extract pixel-wise defect zones from the healthy background, with applicability across all five defined defect categories. Real-time performance and robustness are essential to ensure system scalability in industrial AFP settings.

Prior studies have explored a range of segmentation algorithms, particularly for optical inspection, spanning from simple masking and thresholding to region growing, watershed, and graph-based methods (Kumar (2008); Meister, Wermes, Stüve & Groves (2021)). In this work, we evaluate two complementary algorithms on thermal images: AT and AC. Both leverage texture features previously extracted via GFs, ensuring computational efficiency through feature reuse from earlier detection stages.

As a preprocessing step, the system performs mean image subtraction: the mean thermal profile of healthy lay-ups is computed from the training database and subtracted from incoming defective images. This removes static biases due to inherent material structure or inter-tow gaps, yielding a zero-centered image that enhances contrast between defects and background patterns (see Figure 4.5).

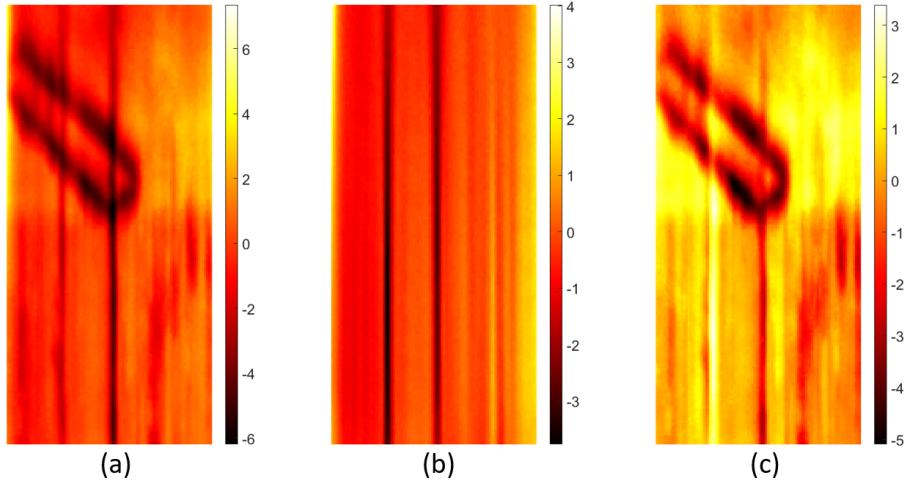


Figure 4.5 Preprocessing via mean subtraction: (a) input image with defect (FOD), (b) mean image of healthy lay-ups, (c) result after subtraction enhancing defect contrast.

The first segmentation method applies Gabor-based enhancement followed by Otsu's adaptive thresholding. While GFs alone provide directional texture discrimination, combining them with AT yields superior segmentation results, as also noted in (Kumar (2008)). Otsu's method assumes a bimodal histogram and computes an optimal

threshold  $T_e$  that maximizes inter-class variance between background and foreground regions (Otsu (1979)). Its real-time applicability is further supported by fast variants tailored for thermal imagery (Bindu & Prasad (2012)).

The second segmentation approach adopts an active contour model, in which an initial curve iteratively evolves to enclose the defective region. Compared to techniques such as standalone Gabor filtering, watershed segmentation, or divide-and-conquer methods, active contours offer superior flexibility and accuracy—especially when properly initialized. Previous research has demonstrated their effectiveness in thermal image segmentation under controlled conditions (Sreeshan, Dinesh & Renji (2020)). Building on this foundation, we extend the method to real-world AFP thermal imagery, where defect shapes are irregular and background noise is nontrivial.

To ensure robust performance, we initialize the contour using texture maps derived from previously extracted GFs rather than random seeds. This strategy improves both localization accuracy and convergence speed. The contour is evolved using the Chan–Vese region-based model (Chan & Vese (2001)), which balances internal smoothness constraints with external forces derived from texture contrast. This energy minimization process enables the contour to adaptively fit the defect boundary, even in soft-gradient thermal contexts. While computationally intensive by design, the use of low-resolution input and precomputed features keeps runtime feasible for near real-time deployment. Upon convergence, the defective area is cleanly isolated from the background and passed to subsequent evaluation modules.

#### 4.4.2 Defect Area Percentage Estimation

In order to transition from visual identification to frame-wise quantitative assessment, a novel metric—DAP—is introduced. DAP is defined as the ratio of defect pixels to the total number of pixels in the affected region along the horizontal lay-up axis. This formulation reflects the anisotropic nature of AFP defects, which primarily impact material integrity transversely.

The segmented defect mask, obtained from either thresholding or active contouring, serves as input to the DAP computation block. Pixel-wise counts are converted into physical units using the thermal camera’s calibration parameters, ensuring dimensional accuracy. Crucially, the evaluation extends beyond frame-level assessment: leveraging the dynamic tow identification mechanism (introduced earlier), DAP is

computed not only for the entire course but also for each individual tow in the frame.

This tow-level granularity enables precise fault localization and supports maintenance strategies aligned with practical AFP workflows, where operators typically intervene at the tow level rather than replacing entire courses. The system computes DAP values for each tow by intersecting the segmented defect mask—obtained via the local inspection framework—with the tow layout map generated by the dynamic tow identification module. Notably, when augmented with the outcomes of the Motion-Aware Global Inspection Framework, this localization gains an additional dimension of continuity. Through inter-frame motion estimation and frame stitching, the second framework provides temporal alignment and spatial consistency across overlapping regions, enabling the traceability of individual defects across the evolving lay-up surface.

The resulting defect metrics are then fused with classifier predictions and relayed to the operator interface alongside intuitive visual overlays. This multi-source integration allows the AFP-DSS to assign each lay-up instance to one of several predefined expert knowledge based quality states, reflecting both local defect severity and broader spatial context. Based on internal thresholds, an alert-recommendation engine issues action prompts ranging from 'No Action Required' to 'Repair or Replace Tow'. Through this hierarchical and context-aware evaluation—anchored by local features and reinforced by global reconstruction—the proposed system elevates thermographic inspection from basic defect detection to an intelligent, operator-assistive quality management platform.

## 5. MACHINE-INDEPENDENT MOTION-AWARE GLOBAL INSPECTION FRAMEWORK

This chapter presents the second core framework of the proposed system: a machine-independent, motion-aware global inspection framework that reconstructs the thermal lay-up from overlapping image sequences. Complementing the local frame-wise defect analysis detailed in Chapter 4, this framework exploits temporal continuity to generate spatially coherent, laminate-scale visualizations. By integrating dense motion estimation with process-aware reconstruction, it enables real-time, machine-independent monitoring and enhances traceability across the evolving lay-up.

### 5.1 Framework Overview

The proposed framework enables real-time, machine-independent inspection of AFP processes by accurately estimating motion fields and reconstructing a spatially consistent global view of the laminate from thermographic video data. It addresses key challenges in motion-aware reconstruction—such as radiometric noise, frame misalignment, tool-dependent variability, and drift accumulation—while emphasizing generality, low latency, and interpretability. Operating independently of machine-integrated signals, the system adapts seamlessly across diverse AFP setups.

As illustrated in Figure 5.1, the system comprises two core stages: **ThermoRAFT-AFP Motion Estimation** and **Thermal Lay-up Reconstruction**. These modules form a sequential pipeline: each frame’s motion is estimated at the pixel level, then aggregated into high-fidelity laminate visualizations. Deep learning is applied selectively—using dense optical flow models where fine-grained temporal consistency is essential, and combining geometric rectification with statistical pooling for efficiency and robustness.

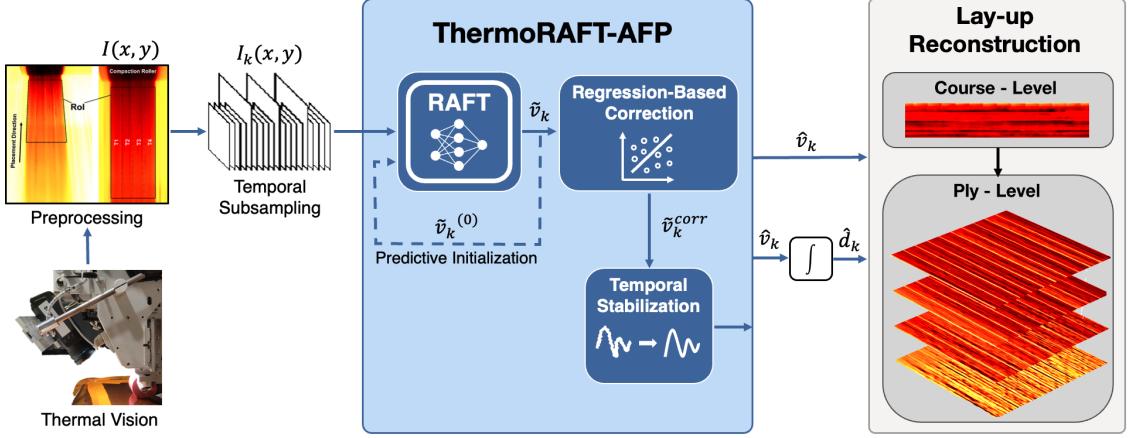


Figure 5.1 Overall structure of the proposed Motion-Aware Global Inspection Framework, illustrating the thermal input stream, motion estimation process, and reconstructed laminate output.

The first module, **ThermoRAFT-AFP Motion Estimation**, processes consecutive thermal frames  $I_k(x, y)$  to compute dense velocity fields  $\hat{\mathbf{v}}_k(x, y)$ :

$$(5.1) \quad \hat{\mathbf{v}}_k(x, y) = \mathcal{F}(I_{k-1}(x, y), I_k(x, y)),$$

where  $\mathcal{F}(\cdot)$  denotes the learned flow estimator. Displacement is recursively accumulated as:

$$(5.2) \quad \hat{\mathbf{d}}_k(x, y) = \hat{\mathbf{d}}_{k-1}(x, y) + \hat{\mathbf{v}}_k(x, y) \cdot \Delta t,$$

with  $\Delta t$  denoting the frame time interval.

To ensure robustness under low-texture and thermally variable conditions, this stage employs a deep learning-based optical flow model with three runtime enhancements: predictive initialization for flow continuity, drift correction via residual alignment, and temporal filtering using exponential moving averages. As a consequence of the perspective correction step in the universal preprocessing pipeline, all input frames are rectified to align the lay-up surface fronto-parallel—simplifying motion estimation to a primarily vertical task. Within this rectified domain, dense vertical flow is pooled into scalar velocity  $v_k$  values—and correspondingly scalar displacement  $d_k$  values—using a median operator over a predefined ROI, yielding compact, drift-stable motion signals well-suited for real-time lay-up reconstruction.

The second module, **Thermal Lay-up Reconstruction**, transforms these scalar displacements into spatially coherent laminate views. Frames are first stitched into seamless course-level mosaics based on vertical motion. These mosaics are then registered ply-wise by matching cumulative displacement profiles. The result is a

temporally ordered, high-fidelity 2D/3D representation of the evolving lay-up surface—enabling traceability, operator guidance, and automated inspection at the ply or laminate scale.

To meet industrial standards and achieve state-of-the-art performance, the system must fulfill:

- **Metrological Performance:** High-precision motion estimation with  $\text{RMSE} < 2\%$  of nominal tool speed,  $R^2 > 99\%$ , and minimal bias ( $\text{MPE} < 1\%$ ).
- **Real-Time Operation:** Minimum 10 fps throughput on standard thermal streams using commodity hardware—without hardware acceleration or synchronization.
- **Deployability:** Robust to thermal noise, invariant to lay-up parameters, and nearly tune-free across diverse process setups—with no machine-derived inputs.
- **Traceability:** Seamless reconstruction with cumulative drift  $\Delta_{\text{drift}} < 1\%$ , enabling high-fidelity 2D mosaics and 3D ply-level visualization for global defect tracking.

Together, these modules compose a unified global inspection framework that complements frame-wise analysis by embedding defect observations within their broader spatiotemporal context. Designed for speed, robustness, and interpretability, the system enables true machine-independent in-situ thermographic inspection—delivering actionable insight into AFP quality evolution.

## 5.2 Deep Flow Estimation in AFP

Deep optical flow models have recently demonstrated superior performance in capturing dense motion patterns, particularly in low-texture or noisy environments. Among them, RAFT (Teed & Deng (2020)) has emerged as a leading architecture, known for its accuracy, resolution preservation, and robustness. RAFT estimates full-resolution optical flow by iteratively refining motion predictions over a high-dimensional all-pairs correlation volume. Its core architecture—depicted in Figure 5.2—comprises a feature encoder, correlation lookup module, and a recurrent update block.

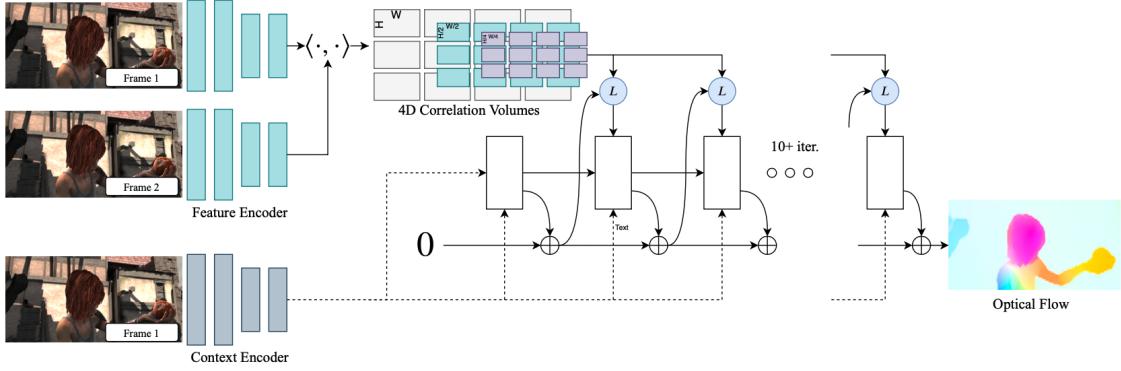


Figure 5.2 RAFT optical flow architecture.

Originally trained on visible-spectrum datasets such as FlyingChairs (Dosovitskiy, Fischer, Ilg, Hausser, Hazirbas, Golkov, Van Der Smagt, Cremers & Brox (2015)), Sintel (Butler, Wulff, Stanley & Black (2012)), and KITTI (Geiger, Lenz, Stiller & Urtasun (2013)), RAFT delivers subpixel-accurate, dense flow across a wide range of real-world scenes. It has achieved top-tier benchmark scores with a relatively modest parameter footprint of 5.3M, as highlighted in recent comparative studies (Alfarano, Maiano, Papa & Amerini (2024b)).

Several architectural properties make RAFT an attractive candidate for thermal motion estimation in AFP settings:

- **Global matching:** All-pairs correlation handles large displacements and sparse features—ideal for untextured thermal inputs.
- **Iterative refinement:** The recurrent update mechanism increases robustness to radiometric fluctuations.
- **Full-resolution output:** Flow maps retain pixel-level precision, which is critical for reconstructing fine-layer AFP structures.

Despite these advantages, raw RAFT is not directly deployable for in-situ AFP monitoring. First, its training on visible-range images renders it insensitive to the unique thermal behaviors of composite materials—such as emissivity variation, heat conduction, and layer stacking effects. Second, its dense matching and iterative inference impose substantial computational load. Even on high-end GPUs, standard RAFT processes  $1088 \times 436$  frames at only  $\sim 10$  fps (Teed & Deng (2020)), falling short of real-time needs. Most importantly, RAFT lacks awareness of AFP-specific motion priors, including vertical tool travel, unidirectional lay-up progression, and temporal coherence.

To address these limitations, we introduce *ThermoRAFT-AFP*: a domain-adapted variant designed for thermal flow estimation in AFP contexts. The model inte-

grates predictive initialization for temporal consistency, radiometric drift correction to handle thermal bias accumulation, and lightweight inference strategies tailored to AFP’s predominantly vertical motion dynamics. Together, these enhancements allow real-time operation while improving accuracy and robustness under true AFP process conditions.

### 5.3 ThermoRAFT-AFP: Domain Tailored Motion Estimator

#### 5.3.1 Predictive Warm-Start Initialization

A central enhancement in our motion estimation strategy is the use of predictive warm-start initialization, which significantly boosts convergence, flow stability, and throughput. Unlike conventional RAFT, which initializes the optical flow field with zeros—implicitly assuming motionless input—or with random values, our approach reuses the final refined output from the previous frame pair as a starting point for the next. This transforms the flow estimator from an open-loop predictor into a temporally guided, fixed-point iterator. For frame  $k$ , the initial flow estimate is defined as:

$$(5.3) \quad \tilde{v}_k^{(0)} = \tilde{v}_{k-1}^{(N)}$$

where  $\tilde{v}_{k-1}^{(N)}$  represents the refined flow at the final iteration  $N$  from the previous pair  $(I_{k-2}, I_{k-1})$ . We set  $N = 12$  consistent with RAFT’s default. This warm-start mechanism reduces refinement steps by 25–40%, enhances stability in low-texture regions, and provides improved flow continuity—thereby enabling real-time operation in process-driven thermal imaging scenarios such as AFP.

#### 5.3.2 Regression-Based Drift Correction

Although RAFT-based models produce high-fidelity motion estimates, small systematic deviations can accumulate over time due to texture sparsity, radiometric

variability, or setup-dependent calibration factors such as lens distortion and pixel-to-millimetre conversion inaccuracies. These per-frame biases, though minor, integrate into significant displacement drift.

To mitigate this, we apply *Passing–Bablok regression* (Passing & Bablok (1983))—a robust, non-parametric linear calibration method designed for uncertain, non-Gaussian data. It requires only a monotonic relationship between predicted and ground-truth values, which we confirm using *Spearman’s rank correlation* (Spearman (1961)). All sequences satisfy the criterion with  $\rho > 0.9$  and  $p < 0.01$ .

Let  $\tilde{v}_k$  denote the raw velocity estimate and  $v_k$  the ground-truth velocity at frame  $k$ . The regression model is given by:

$$(5.4) \quad v_k = \beta \tilde{v}_k + \alpha$$

with slope  $\beta$  and intercept  $\alpha$  estimated from median pairwise slopes. The corrected velocity becomes:

$$(5.5) \quad \tilde{v}_k^{\text{corr}} = \frac{\tilde{v}_k - \alpha}{\beta}$$

This adjustment eliminates systematic drift while preserving signal trends. Importantly,  $(\alpha, \beta)$  are computed once per camera–setup pair using a short calibration sequence and remain fixed thereafter—enabling generalized, nearly calibration-free deployment across varying materials, machines, and lay-up parameters. Despite its simplicity, this correction yields the highest gain in accuracy across our evaluation metrics and incurs negligible runtime cost, making it ideal for embedded inspection workflows.

### 5.3.3 Temporal Stabilization

In addition to accuracy, AFP monitoring demands temporally consistent motion outputs suitable for reconstruction and control. Even with subpixel precision, frame-to-frame jitter may arise due to transient emissivity effects or localized contrast drops. To suppress such instability, we apply a causal exponential moving average (EMA) filter to the corrected velocity stream:

$$(5.6) \quad \hat{v}_k = \lambda \tilde{v}_k^{\text{corr}} + (1 - \lambda) \hat{v}_{k-1}$$

where  $\lambda \in [0, 1]$  is a smoothing gain. This yields the stabilized output  $\hat{v}_k$  with minimal computational burden.

We evaluated alternative smoothing strategies such as the Kalman filter (Kalman (1960)) and Savitzky–Golay filter (Savitzky & Golay (1964)). However, Kalman filtering introduced undesirable parameter sensitivity and minor lag in our setting, while Savitzky–Golay filtering, though shape-preserving, is non-causal and unsuitable for real-time streaming.

EMA, by contrast, provides low-latency smoothing with a single multiply–accumulate operation per frame. It consistently reduces jitter, preserves the underlying trend, and yields small but reliable improvements in estimation quality—completing the *ThermoRAFT-AFP* pipeline with temporally stable, real-time motion output.

## 5.4 Process-Aware Inference Optimizations

Despite the architectural precision and temporal stability of *ThermoRAFT-AFP*, achieving real-time performance in industrial AFP environments requires further optimizations beyond model design. To this end, we introduce a set of domain-specific inference strategies that substantially reduce computational load without compromising accuracy. These are not generic speed-ups, but process-aware accelerations that exploit the predictable motion patterns and structured thermal features inherent to AFP lay-up.

### 5.4.1 Region-of-Interest Cropping

Due to the quadratic scaling of RAFT’s 4D correlation volume (Martins Briedis, Gross & Schroers (2025); Teed & Deng (2020)), reducing spatial resolution directly improves inference speed. In AFP thermography, significant motion is confined to a narrow central band corresponding to the active laminate region, while peripheral zones typically contain static background and radiometric noise.

To exploit this, we define a fixed symmetric ROI per setup, cropping lateral margins

to isolate the motion-rich band. Unlike adaptive or entropy-based cropping, this approach imposes no per-frame overhead and preserves geometric alignment—ensuring compatibility with warm-start initialization and downstream stitching.

Empirically, this ROI gating yields up to 20% speed-up with equal or improved flow stability. By focusing inference on the relevant region, it enhances robustness while maintaining sub-millimetre tracking fidelity at high throughput.

#### 5.4.2 Adaptive Early Exit

RAFT’s iterative refinement allows for early termination once convergence is reached. We exploit this via an adaptive early-exit mechanism that monitors the update norm  $|\Delta\hat{v}|$  at each iteration. Refinement halts when the magnitude falls below a threshold  $\tau$  (empirically  $\tau < 0.01$  px), beyond which further iterations offer negligible gain.

Combined with predictive warm-starts, this strategy enables rapid convergence, often reducing iterations from 12 to as few as 1–3. In effect, it transforms RAFT into a convergence-aware tracker that adapts to the smooth, unidirectional motion typical of AFP processes. This optimization significantly lowers latency while maintaining flow quality.

#### 5.4.3 Temporal Subsampling

RAFT’s robust all-pairs matching allows for inference under sparse temporal sampling, tolerating larger displacements than traditional block- or gradient-based methods. We exploit this via fixed-interval frame skipping, reducing the effective processing rate  $f_{\text{process}}$  below the camera’s nominal rate  $f_{\text{nom}} = 50$  Hz.

To regulate this trade-off, we define the inter-frame displacement  $\delta$  and the vertical overlap ratio  $\gamma$  as:

$$(5.7) \quad \delta = v_{\text{nom}} \cdot \frac{\rho}{f_{\text{process}}}$$

$$(5.8) \quad \gamma = 100 \cdot \left( 1 - \frac{\delta}{h_{\text{RoI}}} \right)$$

where  $v_{\text{nom}}$  is the nominal tool speed (mm/s),  $\rho$  is pixel density (px/mm), and  $h_{\text{RoI}}$  is the vertical height of the RoI (px).

As an example, using  $v_{\text{nom}} = 250 \text{ mm/s}$ ,  $\rho = 4 \text{ px/mm}$ , and  $h_{\text{RoI}} = 180 \text{ px}$ , selecting  $f_{\text{process}} = 25 \text{ Hz}$  (i.e., `StepSize = 2`) yields  $\delta \approx 40 \text{ px}$  and  $\gamma \approx 77.8\%$ . This provides a smooth overlap for accurate stitching while halving computational cost.

Unlike entropy- or event-based policies, this deterministic approach is simple, robust, and compatible with defect detection and reconstruction. Among all proposed optimizations, temporal subsampling yields the largest runtime savings—making it critical for real-time deployment.

## 5.5 Reconstructing the AFP Process

While frame-wise velocity estimation yields temporally localized insights, industrial AFP monitoring increasingly requires laminate-scale visualizations to support defect traceability, maintenance diagnostics, and intelligent process control. Operators typically assess plies globally but intervene at the course or tow level—necessitating both coarse and fine-grained visual context.

To mirror this hierarchical inspection logic, our system autonomously reconstructs the lay-up at two levels. First, thermal frames are stitched into seamless course-level mosaics. Then, these mosaics are vertically aligned into ply-scale assemblies using inter-course velocity synchronization. This design supports operator interpretation, enables automated evaluation, and ensures spatiotemporal coherence. The full pipeline is summarized in Algorithm 2.

### 5.5.1 Course-Level Seamless Image Stitching

Each course  $C_i$  consists of a sequence of thermal frames  $I_k$ , for which corrected velocity  $\hat{v}_k$  and displacement  $\hat{d}_k$  are available from *ThermoRAFT-AFP*. Frames are vertically aligned using  $\hat{d}_k$  and blended into the course mosaic  $\mathcal{C}_i$  using a NaN-aware

---

**Algorithm 2** Real-time multi-level AFP lay-up reconstruction algorithm

---

**Input:** Frame  $I_k$ , velocity  $\hat{v}_k$ , displacement  $\hat{d}_k$ , timestep  $\Delta t$ , density  $\rho$   
**Output:** Reconstructed laminate  $\mathcal{L}$

$i \leftarrow 1$ ,  $\mathcal{L} \leftarrow \emptyset$ ,  $\mathcal{P}_\ell \leftarrow \emptyset$

**while** receiving frame  $I_k$  **do**

**if** first frame of course  $C_i$  **then**

$\mathcal{C}_i \leftarrow I_k$ ,  $V_i \leftarrow [\hat{v}_k]$

**else**

Shift  $I_k$  by  $\hat{d}_k$ , blend into  $\mathcal{C}_i$  (Eq. 5.9)

Append  $\hat{v}_k$  to  $V_i$

**if** motion stop detected **then**

**if**  $i = 1$  **then**

$\mathcal{P}_\ell \leftarrow \mathcal{C}_1$

**else**

$\ell_i \leftarrow \arg \max_\ell \text{xcorr}(V_i, V_{i-1})$

$s_i \leftarrow \left( \sum_{j=1}^{|\ell_i|} |\hat{v}_i[j]| \cdot \Delta t \right) \cdot \rho$

Vertically align  $\mathcal{C}_i$  by  $s_i$ , append to  $\mathcal{P}_\ell$

$i \leftarrow i + 1$

Append  $\mathcal{P}_\ell$  to  $\mathcal{L}$ , **return**  $\mathcal{L}$

---

feathering scheme to ensure radiometric continuity and eliminate edge artifacts. The blended pixel intensity  $B(x, y)$  is computed as:

$$(5.9) \quad B(x, y) = \frac{\sum_k w_k(x, y) \cdot I_k(x, y) \cdot \delta_k(x, y)}{\sum_k w_k(x, y) \cdot \delta_k(x, y)}$$

where  $w_k(x, y)$  is a tapered spatial weight (e.g., linear or Gaussian), and  $\delta_k(x, y)$  is the binary validity mask:

$$(5.10) \quad \delta_k(x, y) = \begin{cases} 1 & \text{if } I_k(x, y) \text{ is valid (non-NaN)} \\ 0 & \text{otherwise} \end{cases}$$

This ensures seamless merging of overlapping regions while excluding missing or corrupted pixels. Each new course is initialized from its first valid frame, and segmentation is driven entirely by motion cues. Specifically, thresholded plateaus of near-zero velocity—indicative of roller deceleration and lift-off—mark the start or end of a course. This enables fully autonomous, data-driven segmentation without reliance on machine-side triggers or external timestamps.

### 5.5.2 Laminate-Level Assembly via Motion Synchronization

Once all courses  $\{\mathcal{C}_i\}$  within a ply are stitched, they are vertically registered into a complete ply view. This is accomplished via motion synchronization: each course's velocity trace  $V_i$  is cross-correlated with that of its predecessor to find the optimal vertical offset:

$$(5.11) \quad s_i = \arg \max_{\ell} \text{corr}(V_i(t), V_{i-1}(t + \ell))$$

This shift  $s_i$  is then used to vertically align  $\mathcal{C}_i$  with  $\mathcal{C}_{i-1}$ . The resulting ply-level mosaics are stacked in build direction, yielding a 3D thermal laminate  $\mathcal{L}$  that is both spatially and temporally consistent. This layered reconstruction enables part-scale analysis, multi-defect traceability, and operator-level interpretation.

Critically, the result serves as a digital substrate for closed-loop AFP control, historical diagnostics, and intelligent repair guidance. By elevating frame-wise observations into a globally coherent visual context, this module transforms inspection from passive defect reporting to active process understanding.

## 6. EXPERIMENTAL RESULTS AND DISCUSSION

This chapter presents a comprehensive evaluation of the hybrid inspection system, structured around the dual-framework architecture from Chapters 4 and 5. Each framework is assessed using thermal datasets totaling over 13,000 frames and evaluated against the success criteria in Section 1.4, covering accuracy, runtime, robustness, and deployability. Results are presented by framework: frame-wise defect analysis and quality assessment are followed by motion estimation and laminate reconstruction. The chapter concludes with a benchmark comparison against state-of-the-art AFP inspection methods.

### 6.1 Multi-Resolution Thermal Dataset for Dual Framework Evaluation

The evaluation of both proposed frameworks was conducted using thermal image data acquired and preprocessed with the setup described in Chapter 3. Due to the lack of publicly available thermal datasets for AFP processes, all evaluation data was acquired in-house. Two task-specific evaluation datasets were then constructed from this data, each reflecting a distinct temporal sampling strategy. The first, used for frame-wise defect analysis and quality assessment (Framework 1), consists of non-overlapping frames selected to maximize information diversity and minimize redundancy. The second, used for motion-aware reconstruction (Framework 2), retains the full sequential frame order, as temporal overlap is essential for accurate motion estimation and stitching.

### 6.1.1 Framework 1 Evaluation Dataset:

A curated subset of 5,000 thermal images was extracted to evaluate detection, classification, and segmentation tasks. This dataset includes 2,500 healthy lay-up frames and 2,500 defective frames, evenly distributed across five common AFP defect categories: gaps, overlaps, missing tows, tow splices, and foreign object debris (FOD). Each image was annotated with both class labels and pixel-wise segmentation masks. The dataset supports module-level evaluation of binary health classification, multi-class defect recognition, and DAP-based segmentation performance.

### 6.1.2 Framework 2 Evaluation Dataset:

The full thermal image sequence was used to evaluate dense motion estimation and lay-up reconstruction. This dataset comprises 13,300 sequential frames captured across 70 AFP courses. Approximately 30% of the sequences contain induced defects, while the remainder represent healthy or low-contrast lay-ups. These sequences were used to compute displacement profiles, stitching drift, and velocity error metrics, forming the quantitative basis for evaluating Framework 2's global inspection performance.

Together, these two purpose-aligned datasets enable rigorous, resolution-specific assessment of both frameworks within the proposed hybrid system.

## 6.2 Framework 1 Evaluation: Defect Analysis and Quality Assessment

This section provides a focused evaluation of the first proposed framework, emphasizing its defect detection, classification, and localization capabilities. It incorporates evaluation results, figures, and tables adapted from the author's previously published works (Zemzemoglu & Unel (2023,2); Zemzemoglu et al. (2024)), in accordance with the respective publishers' reuse policies. The analysis is organized by subsystem: dynamic tow identification, hierarchical defect identification, and defect segmentation. For each module, we examine performance using both quantitative metrics and illustrative examples, with special attention to robustness under defect-induced noise.

Runtime results confirm the system's adherence to real-time constraints, including under full execution scenarios. The findings affirm the framework's suitability for adaptive, frame-wise AFP inspection in industrial settings.

### 6.2.1 Performance Analysis of Dynamic Tow Identification Algorithm

We evaluate the performance of the proposed spatial-temporal analysis algorithm for dynamic tow identification in AFP lay-up images. To underscore the contribution of the temporal component, we begin by examining spatial-only estimations. Figure 6.1 shows results on one healthy and one defective sample, comparing estimated (blue dashed) and ground-truth (green solid) tow boundaries. The accompanying edge histograms (*vHist*) help visualize gradient distributions critical for boundary inference.

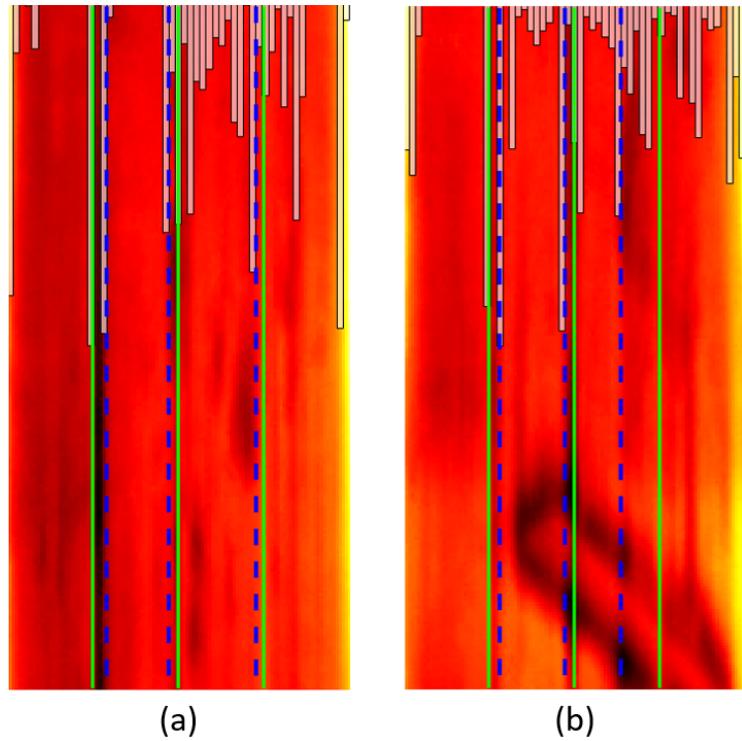


Figure 6.1 Tow boundary estimation using spatial-only analysis on (a) healthy and (b) defective samples with ground truth (green lines) and estimated boundaries (blue dashed lines); bar charts show edge distributions.

The calculated pixel-wise error margins for this spatial tow boundary estimation example are  $[+4.5, -3, -2.5]$  for the healthy lay-up and  $[+3.5, -3, -12.5]$  for the defective lay-up. While the spatial algorithm performs reasonably on healthy samples, defective regions—such as overlaps or tow splices—induce substantial de-

viations that impair tow-level reasoning, particularly DAP computation. Despite parameter tuning (e.g., threshold  $T$ , merging distance  $mDist$ ), spatial-only methods lack the robustness required under defect-induced noise. These limitations motivate our integration of a temporal component, which stabilizes estimations with minimal additional computational cost.

A detailed scenario is presented in Figure 6.2, where two defect events—D1 (tow splice) and D2 (overlap)—emerge during the lay-up. The splice in tow 3 introduces outlier vertical edges, triggering a spike in B3’s estimation error and momentarily pushing B2 outside the convergence band. The subsequent overlap between tows 3 and 4 causes temporary divergence in B3 before recovery. Despite these disturbances, the spatial-temporal algorithm adapts in real time, progressively refining estimates through cumulative averaging. Both B2 and B3 ultimately converge within the  $\pm 2$  px error margin, demonstrating the algorithm’s resilience and alignment with the robustness criterion defined for tow identification.

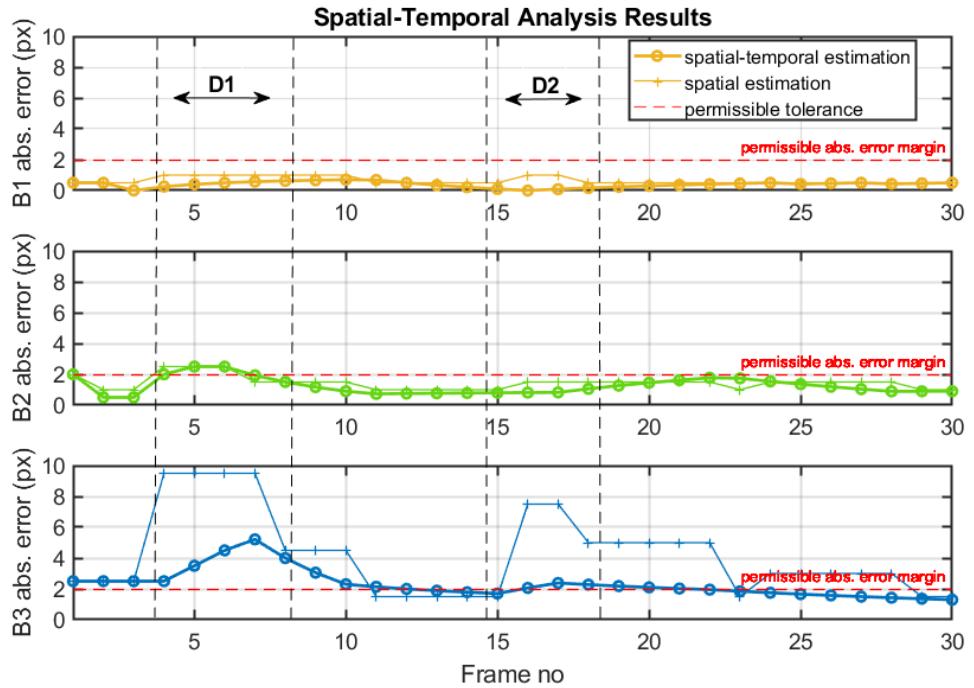


Figure 6.2 Tow boundary errors under defects D1 and D2. Temporal averaging stabilizes B2 and B3 within  $\pm 2$  px band.

To quantify performance, we define four metrics: (i) *First Convergence*, the earliest frame index where error enters the tolerance envelope; (ii) *Stable Convergence*, where it remains consistently within bounds; (iii) *Estimation Error*, the mean absolute boundary error; and (iv) *Execution Time*, measuring runtime per frame. These metrics directly reflect the system’s **Tow Identification Robustness** and contribute to its overall real-time suitability.

The algorithm parameters were experimentally optimized for balanced accuracy, robustness, and runtime. Key settings included a peak merging distance  $mDist = 10$  px, edge threshold  $T = 0.5$ , and a 10-frame temporal window ( $k = 10$ ) for averaging. The weighting array  $w$  was designed to favor boundary zones over the center, mitigating the impact of defective regions. Table 6.1 compares the performance of different filter-averaging combinations across four criteria.

Among the variants, Scharr-based linear filtering with cumulative averaging ( $\mathbf{C}_{avg}$ ) consistently achieved the best trade-off—delivering the lowest estimation error (0.8 px), fastest stable convergence (8 frames), and sub-30 ms runtime. This performance reflects the algorithm’s momentum-driven convergence: as the majority of frames are defect-free, cumulative averaging increasingly suppresses defect-induced noise. In contrast, non-linear filters (e.g., top-hat) yielded no significant gain but added computational cost. Overall, the selected configuration meets the real-time requirement and aligns with the system’s **Tow Identification Robustness** criterion.

Table 6.1 Mean performance comparison of the proposed spatial-temporal analysis based dynamic tow identification algorithm using different filtering methods and averaging techniques.

	No Filter			Linear Filter			Non-linear Filter		
	SM <sub>avg</sub>	WM <sub>avg</sub>	C <sub>avg</sub>	SM <sub>avg</sub>	WM <sub>avg</sub>	C <sub>avg</sub>	SM <sub>avg</sub>	WM <sub>avg</sub>	C <sub>avg</sub>
First Convergence [f]	8.2	7.0	7.4	4.2	4.0	<b>3.4</b>	5.9	5.5	4.7
Stable Convergence [f]	14.5	14.0	12.2	9.5	8.9	<b>8.0</b>	12.0	10.8	10.5
Estimation Error [px]	3.6	3.5	3.1	1.3	1.1	<b>0.8</b>	2.2	1.8	1.8
Execution Time [ms]	21.7	23.3	22.0	26.1	28.9	<b>27.7</b>	33.5	35.1	31.6

The chosen configuration balances robustness and efficiency, with convergence typically achieved in under 8 frames and sub-pixel accuracy maintained across healthy and defective regions. Compared to nonlinear alternatives, linear filtering provides competitive performance at significantly lower computational cost.

To further enhance estimation reliability, classification outputs from the defect detection module can optionally guide dynamic re-weighting or masking strategies—skipping highly defective frames or down-weighting their influence in the averaging process. This integration could improve robustness against emerging defects while maintaining low overhead.

Unlike prior works that rely on static assumptions or machine-coupled models, our method operates fully online, adapting to new frame content and enabling generalized deployment across AFP systems. This approach aligns directly with the **Tow Identification Robustness** and **Real-Time Operation** criteria outlined earlier.

### 6.2.2 Hierarchical Defect Identification Performance Evaluation

This part evaluates the hierarchical learning-based framework for frame-wise defect identification. The system consists of a high-level binary classifier using Support Vector Machines (SVM) and a low-level multiclass classifier based on deep convolutional neural networks (DCNN). Performance is reported in accordance with the defined criteria for **Detection Reliability** and **Classification Accuracy**, with attention to recall, false negatives, and runtime feasibility.

#### 6.2.2.1 High-Level SVM-Based Defect Detector

The SVM model was trained using Gabor texture features extracted from the labeled thermal dataset. The data was partitioned into three subsets:

- **Training set (80%)**: Used to fit the model.
- **Validation set (10%)**: Used to optimize hyperparameters during training.
- **Test set (10%)**: Held out for final evaluation using unseen instances.

A grid search was employed to identify optimal hyperparameters, with tuning prioritized toward high recall (sensitivity) to minimize the likelihood of undetected defects. No signs of overfitting were observed, and all reported metrics refer to the test set performance.

Figure 6.3 presents the normalized confusion matrix, treating defective lay-ups as the positive class. The matrix displays counts of true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN) for binary classification. The model achieved 97.2% recall for defective and 95.6% precision for healthy samples. The 2.8% false negative rate ensures most defective instances are flagged, while the 4.4% false positive rate minimizes unnecessary process interruptions and associated costs.

Table 6.2 summarizes the tuned parameters and test metrics. The model achieved an overall accuracy of 96.4% and an F1-score of 96.43%, underscoring its balanced performance. Some misclassifications stem from ambiguities between defect and non-defect textures (e.g., gaps vs. tow boundaries), which may be reduced with conditional postprocessing logic.

Feature extraction dominates the model's average runtime (89.7 ms), mainly due

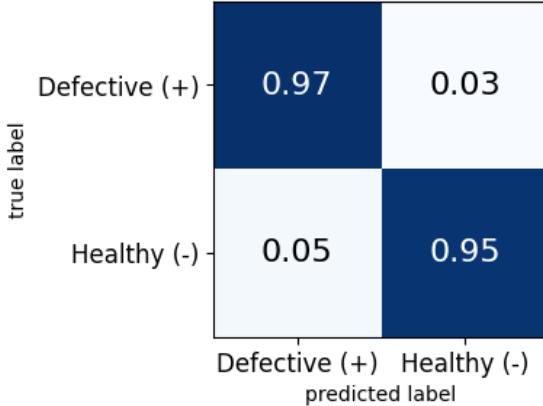


Figure 6.3 Normalized confusion matrix for SVM-based test set classification results.

Table 6.2 SVM model optimized parameters and performance evaluation.

Model Parameters		Performance Metrics	
Kernel constant ( $\gamma$ )	$2^{-5}$	Test accuracy	96.4%
Penalty term ( $C$ )	$2^{-3}$	F1-score	96.43%
Feature extraction	72.7 ms	Recall	97.2%
SVM prediction	17.0 ms	Precision	95.67%
Total time	89.7 ms		

to image convolution with the Gabor filter bank. Although the filters are pre-generated and reused across frames and quality modules, this step remains the most demanding. Still, execution time remains well within the **Real-Time Operation** constraint, especially since the extracted features are reused for downstream quality evaluation. This efficiency, paired with strong predictive performance, confirms the SVM model's suitability for reliable AFP status monitoring.

### 6.2.2.2 Low-Level DCNN-Based Multiclass Classifier

The DCNN model was trained to classify five defect types. Training converged after 60 epochs, with the learning rate reduced every 10 epochs. Key parameters and results are listed in Table 6.3.

Figure 6.4 shows the normalized confusion matrix. All per-class accuracies exceeded 95%. The *Overlap* class posed the greatest challenge, likely due to constraints in synthetic defect generation. *Missing Tow*, despite its low visual contrast and difficulty for human inspection—especially in unidirectional multi-layer lay-ups—achieved the highest accuracy (98%), aided by multi-tow scenarios that improved class separa-

Table 6.3 DCNN model parameters and classification results.

Model Parameters		Classification Results		
Layers number	12	Overall accuracy	96.4%	
Initial learning rate	$10^{-3}$	Test time		39.7 ms
Epochs limit	60			
Training set	80%			
Validation set	10%			
Test set	10%			

bility. *Tow Splice* and *Foreign Body* classes exhibited slightly higher false positives due to inter-class visual similarity.

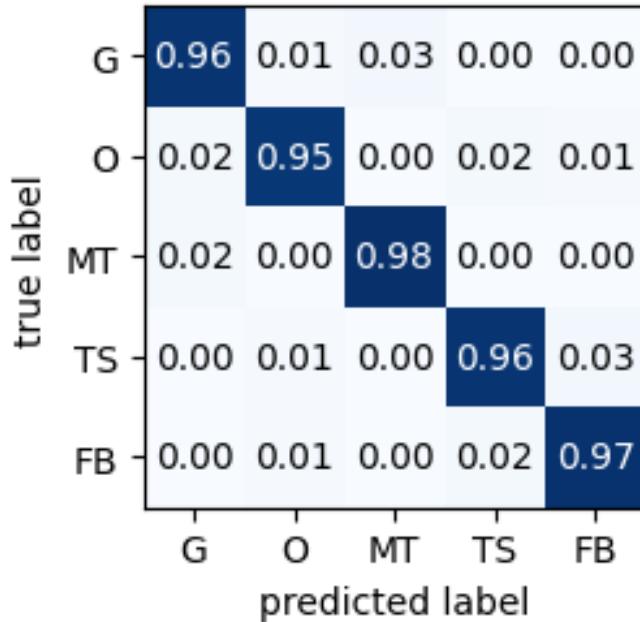


Figure 6.4 Normalized confusion matrix for DCNN-based test set classification results.

Despite minor confusion, all class accuracies exceeded 96%, validating the DCNN’s effectiveness and satisfying the **Classification Accuracy** criterion. The hierarchical structure mitigates class imbalance and simplifies model complexity, enabling more stable generalization. Considering that preprocessing consumes an average of 1.5 ms, the combined detection (89.7 ms) and classification (39.7 ms) stages yield a total module runtime of 130.9 ms per input frame—well within the 5 fps real-time constraint and leaving sufficient margin for the upcoming segmentation stage.

This architecture enables high-confidence, efficient classification while complementing the detection stage to strengthen overall frame-wise defect intelligence. Together, these modules validate the system’s capability for robust, real-time defect detection

and classification—fulfilling the industrial demands of quality-focused AFP monitoring.

### 6.2.3 Defect Segmentation and Localization

The effectiveness of the defect segmentation and localization algorithm is critical to downstream quality evaluation, particularly for computing course-level and tow-level DAP metrics. Two algorithmic variants are compared: adaptive thresholding (AT) of texture features and AC with texture feature seeding. Among defect types, foreign bodies and tow splices pose segmentation challenges due to their irregular shapes, non-uniform thermal contrast, and high similarity to structural textures in the background.

Figure 6.5 shows the result of the AT-based method on a sample image containing a foreign body defect spanning the third and fourth tows. While the texture-based preprocessing enhances defect edges, it also amplifies weaker gradients near tow boundaries and image margins—particularly the middle tow boundary with a wide footprint. These outliers propagate through thresholding and appear as false positives in the final binary mask. Although the sides of the defect are partially segmented, the AT result fails to form a closed region and includes residual pixels from background structures.

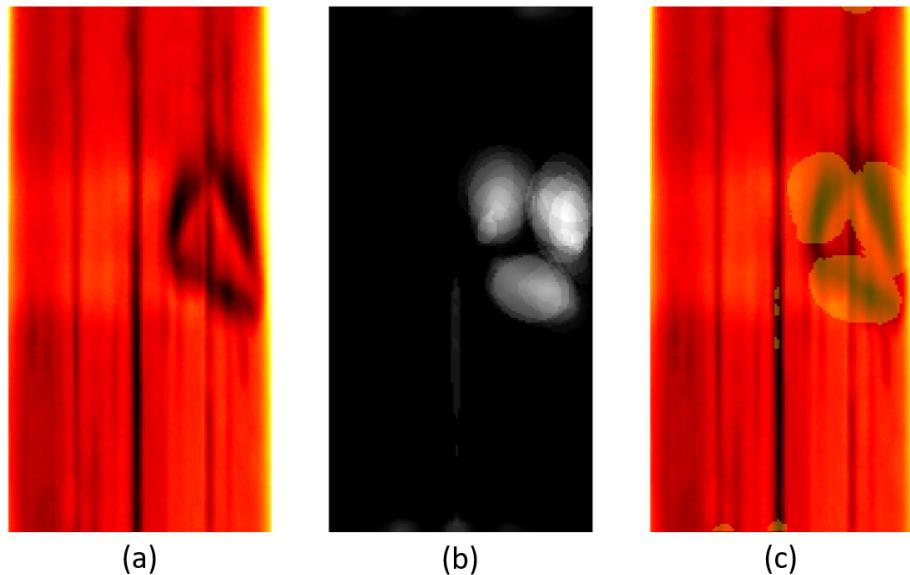


Figure 6.5 Adaptive thresholding of texture features: (a) original lay-up image showing a tow scrap foreign body defect; (b) texture feature-enhanced image; (c) segmented mask using Otsu's adaptive thresholding.

The AC algorithm was applied to the same image for comparative evaluation. As shown in Figure 6.6, the initial seed contains edge clusters from the defect as well as outlier pixels on the middle tow boundary. Over successive iterations, the active contour expands across the defect area while suppressing the influence of non-defective textures—demonstrating strong noise rejection. By iteration  $i = 25$ , the contour converges on the foreign body region, successfully enclosing it. Although the final mask includes some peripheral pixels not part of the true defect, this conservative tendency introduces a slight positive bias in the DAP metric—considered preferable to the risk of false negatives.

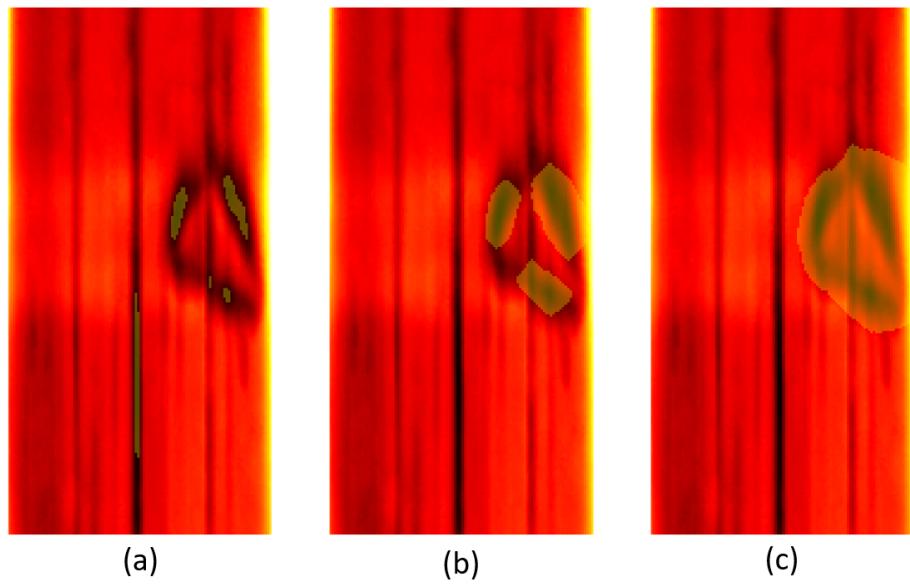


Figure 6.6 Active contours with texture feature seeding: (a) initial seed mask at  $i = 1$ ; (b) intermediate contour at  $i = 10$ ; (c) final convergence at  $i = 25$ , isolating the defect.

Quantitatively, AT excels in computational efficiency with an average runtime of 8.2 ms. The AC algorithm, requiring 62.5 ms for 25 iterations, is slower but offers significantly higher spatial precision and robustness. Its ability to form closed regions and resist false segmentation makes it suitable for integration into real-time quality inspection, especially when single-defect instances are expected. With future integration of region splitting and merging logic, the method can support multi-defect scenarios.

To evaluate accuracy, we adopt two common metrics: PA and IoU. PA measures the proportion of correctly classified pixels, while IoU quantifies the overlap between predicted and ground-truth regions:

$$(6.1) \quad \text{PA} = \left( \frac{TP + TN}{TP + TN + FP + FN} \right) \times 100$$

$$(6.2) \quad \text{IoU} = \frac{|S \cap G|}{|S \cup G|}$$

Here,  $TP$ ,  $TN$ ,  $FP$ , and  $FN$  denote true positives, true negatives, false positives, and false negatives, respectively. In the IoU expression,  $S$  is the predicted segmentation and  $G$  is the ground truth.

The AC method was tested on defective lay-up instances of the evaluation dataset, yielding a mean PA of 93.2% and mean IoU of 0.72. These results highlight the algorithm’s suitability for thermography-based segmentation tasks and confirm its general robustness. Direct comparison with external benchmarks is limited by the absence of publicly available AFP thermographic datasets. However, similar methods operating on profilometer data or bounding-box annotations report IoUs between 0.70–0.78 (Ghamisi et al. (2023); Tang et al. (2022)), suggesting our performance is on par or better given the greater complexity of thermal data.

Taken together, this segmentation strategy provides a reliable defect localization backbone for frame-wise AFP quality analysis. While AT offers a fast but coarse approximation, AC delivers higher accuracy, especially in cases involving irregular or low-contrast defects—ensuring downstream DAP computation remains robust and interpretable.

**Runtime Summary:** The total execution time of the frame-wise quality assessment framework depends on the active route. Route A—comprising only tow identification—was previously shown to complete within a sub-30 ms runtime. In the worst-case Route B scenario—where both classification and segmentation via AC are triggered upon defect detection—the runtime reaches 193.4 ms. This remains within the system’s 200 ms frame budget, satisfying the 5 FPS real-time constraint and confirming the deployability of the full analysis path.

### 6.3 Framework 2 Evaluation: Motion-Aware Global Inspection

This section evaluates the second framework of the proposed inspection system, which leverages overlapping thermal frames to enable motion-aware analysis and global lay-up reconstruction. Emphasis is placed on subpixel motion estimation, real-time feasibility, and laminate-scale visualization. The assessment spans accuracy, drift, and robustness—supporting the framework’s role as a globally consistent

and operator-assistive inspection layer.

### 6.3.1 ThermoRAFT-AFP Performance Analysis

This subsection presents a quantitative evaluation of the ThermoRAFT-AFP framework on a representative thermal dataset comprising 13,300 frames captured from AFP lay-up operations. The evaluation emphasizes subpixel motion estimation accuracy, long-term displacement stability, and temporal consistency. Three metric categories are used: **(i) Precision**—quantified by Root Mean Square Error (RMSE) and Mean Absolute Error (MAE); **(ii) Trueness**—captured by Mean Prediction Error (MPE); and **(iii) Tracking Fidelity**—measured using Coefficient of Determination ( $R^2$ ), Pearson Correlation ( $r$ ), and cumulative drift ( $\Delta_{\text{drift}}$ ). Given the ground truth velocity  $v_k$  and predicted velocity  $\hat{v}_k$  at frame  $k$ , with cumulative displacement  $\hat{d}_k$ , the evaluation metrics are defined as:

$$(6.3a) \quad \text{RMSE} = \sqrt{\frac{1}{N} \sum_{k=1}^N (v_k - \hat{v}_k)^2}$$

$$(6.3b) \quad \text{MAE} = \frac{1}{N} \sum_{k=1}^N |v_k - \hat{v}_k|$$

$$(6.3c) \quad \text{MPE} = 100 \cdot \frac{1}{N} \sum_{k=1}^N \frac{\hat{v}_k - v_k}{v_k}$$

$$(6.3d) \quad R^2 = 1 - \frac{\sum_k (v_k - \hat{v}_k)^2}{\sum_k (v_k - \bar{v})^2}$$

$$(6.3e) \quad r = \frac{\sum_k (v_k - \bar{v})(\hat{v}_k - \bar{\hat{v}})}{\sqrt{\sum_k (v_k - \bar{v})^2} \sqrt{\sum_k (\hat{v}_k - \bar{\hat{v}})^2}}$$

$$(6.3f) \quad \Delta_{\text{drift}} = 100 \cdot \left| \frac{\hat{d}_N - d_N}{d_N} \right|$$

A grid search was conducted on a validation subset to jointly tune the Region-of-Interest (RoI) cropping ratio, RAFT iteration count, and Exponential Moving Average (EMA) gain  $\lambda$ . The selected configuration—20% cropping, two iterations,  $\lambda = 0.82$ , and Passing–Bablok regression correction ( $\beta = 0.9789$ ,  $\alpha = 0.0095$ )—achieves real-time operation while ensuring high precision.

A sample course result is illustrated in Figure 6.7, comparing ThermoRAFT-AFP

with raw RAFT and ground truth. The reference velocity was interpolated using shape-preserving cubic splines to replicate typical lay-up profiles featuring velocity ramps and plateaus. While raw RAFT exhibits overestimation during steady motion, ThermoRAFT-AFP shows tight alignment, particularly in transient zones. Raw RAFT yields 7.87 mm/s RMSE, 3.64% MPE, and 2.58% displacement drift, whereas ThermoRAFT-AFP reduces these to 4.58 mm/s, 0.17%, and 0.10% respectively—achieving sub-millimeter cumulative error over a 750 mm lay-up.

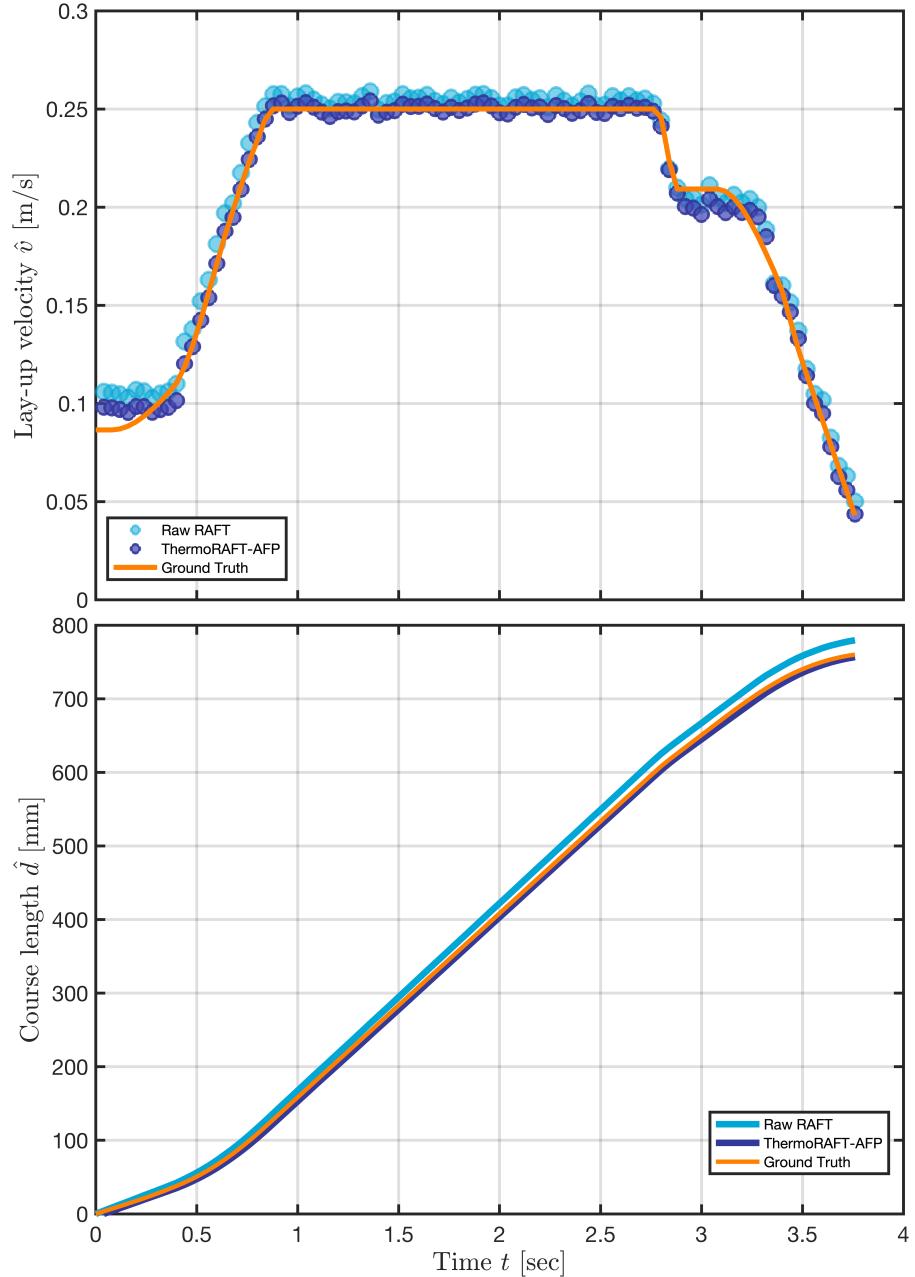


Figure 6.7 Motion estimation for a representative AFP course: (top) frame-wise lay-up velocity, (bottom) cumulative displacement.

The refined estimates track both amplitude and phase accurately, with  $R^2 = 99.51\%$  and Pearson  $r = 99.76\%$ , underscoring fidelity even during dynamic transitions. This

single-case performance typifies the global dataset-level findings reported next.

Table 6.4 summarizes ThermoRAFT-AFP’s metrics across the full evaluation set. Average RMSE remains under 5 mm/s, MPE under 0.2%, and cumulative drift at 0.09%—all satisfying the predefined success thresholds. Compared to baseline RAFT, RMSE and MAE improve by over 40%, while MPE and drift reduce by more than 90%. Moreover, standard deviations are notably reduced, indicating robustness across lay-up scenarios, thermal patterns, and defect contexts. This consistency is further evidenced by high  $R^2$  (99.33%) and correlation (99.70%).

Table 6.4 Performance metrics of ThermoRAFT-AFP across evaluation dataset

	<b>RMSE</b> [mm/s]	<b>MAE</b> [mm/s]	<b>MPE</b> [%]	<b><math>R^2</math></b> [%]	<b><math>r</math></b> [%]	<b><math>\Delta_{\text{drift}}</math></b> [%]
Raw RAFT	$8.70 \pm 1.12$	$6.39 \pm 0.92$	$3.06 \pm 0.58$	$96.7 \pm 0.22$	$98.3 \pm 0.17$	$2.58 \pm 1.96$
ThermoRAFT-AFP	<b><math>4.83 \pm 0.25</math></b>	<b><math>3.70 \pm 0.19</math></b>	<b><math>0.17 \pm 0.25</math></b>	<b><math>99.33 \pm 0.07</math></b>	<b><math>99.70 \pm 0.10</math></b>	<b><math>0.09 \pm 0.08</math></b>
Enhancement <sup>†</sup>	44.5%	42.1%	94.4%	+2.6 pp	+1.4 pp	96.5%
Success criteria	<5.0	<5.0	<0.5	>99.0	>99.0	<1.0

<sup>†</sup>Relative gain (baseline – ours)/baseline × 100; “pp” = percentage points.

Most of the drift and bias correction ( 90%) stems from the regression layer, while EMA filtering contributes an additional 7–10% accuracy gain by mitigating jitter. The integrated enhancements transform a generic flow estimator into a domain-adapted, precision-grade tracker suitable for automated thermographic monitoring in AFP.

### 6.3.2 Runtime Performance and Real-Time Feasibility

This subsection evaluates the real-time operability of the ThermoRAFT-AFP framework by analyzing the trade-off between inference speed and estimation fidelity. The analysis is supported by both metric-based profiling and the effect of cumulative optimization strategies. All results are measured on the hardware platform previously described in chapter 3. Root Mean Square Error (RMSE) is selected as the primary representative metric, while correlation-based indicators such as  $R^2$  and  $r$  are omitted here, having already been validated in subsection 6.3.1.

Table 6.5 summarizes the effect of temporal subsampling (StepSize variation) on both accuracy and latency. The tuned configuration, StepSize 2 (25 Hz), emerges as Pareto-optimal—achieving subpixel velocity error, a low drift rate of 0.09%, and real-time throughput with a measured inference time of 33.0 ms. This is comfortably

within its 40 ms frame cycle budget. Interestingly, StepSize 1 (50 Hz) performs worse despite higher sampling density. Its 70 ms latency exceeds the allowable time and yields degraded accuracy, likely due to excessive spatial redundancy and radiometric noise, which compromise RAFT convergence.

Table 6.5 Accuracy vs inference time trade-off analysis for different StepSize settings.

Metric	StepSize 1	StepSize 2	StepSize 3	StepSize 4	StepSize 5	StepSize 6
	(50 Hz)	(25 Hz)	(16.6 Hz)	(12.5 Hz)	(10.0 Hz)	(8.33 Hz)
RMSE [mm/s]	6.23	<b>4.83</b>	7.18	15.35	10.11	75.98
MPE [%]	-0.39	<b>0.17</b>	1.50	2.66	0.17	-6.62
$R^2$ [%]	99.05	<b>99.33</b>	98.79	94.46	97.22	-54.81
$\Delta_{\text{drift}}$ [%]	0.28	<b>0.09</b>	0.41	0.30	0.35	7.41
Allowable Time [ms]	20	40	60	80	100	120
Inference Time [ms]	70.0	<b>33.0</b>	24.3	18.1	14.6	12.4

The performance degrades sharply at StepSize 6 (8.33 Hz), where RMSE exceeds 75 mm/s and the coefficient of determination ( $R^2$ ) falls below zero. This collapse stems from insufficient overlap between consecutive frames, violating optical flow assumptions and impeding convergence. Intermediate configurations (e.g., StepSize 4 at 12.5 Hz) already exhibit instability, suggesting that temporal motion continuity becomes unreliable below a practical threshold of 16 Hz. These results confirm that high frame rate alone is not sufficient—temporal coherence is critical to accurate flow estimation.

While inference time decreases monotonically with increasing StepSize, accuracy does not degrade linearly. For instance, StepSize 3 is 26% faster than StepSize 2, yet it suffers from 48% higher RMSE and over four times more drift. StepSize 5 maintains reasonable accuracy with 10 Hz processing and may be useful under fallback or power-saving conditions. The selected StepSize 2 configuration also leaves a 7 ms budget margin for preprocessing and reconstruction tasks, validating the feasibility of the full end-to-end pipeline.

ThermoRAFT-AFP achieves this real-time performance not by simplification but through careful domain-specific optimization of the inference pipeline. Starting from baseline RAFT inference—78.5 ms per frame with full resolution and 12 iterations—a series of runtime optimizations are applied. Predictive warm-starting reduces this by 23%, reaching 60.6 ms. ROI cropping across horizontal and vertical axes yields a further 24% improvement, lowering latency to 45.8 ms. Finally, a confidence-driven early-exit mechanism dynamically trims redundant refinement iterations, arriving at the deployed runtime of 33.0 ms.

This 58% cumulative latency reduction demonstrates how domain-aware strategies—not generic simplification—enable real-time inference with high accuracy. ThermoRAFT-AFP thus bridges the gap between academic optical flow models and the practical demands of autonomous, in-situ AFP monitoring and inspection.

### 6.3.3 Lay-up Reconstruction and Visualization

Building on the validated velocity estimates, this subsection demonstrates how ThermoRAFT-AFP enables real-time laminate-scale reconstruction using only thermal imagery. The selected case features a uni-directional (UD) laminate comprising both healthy zones and FOD anomalies.

Figure 6.8 illustrates the course-level reconstruction. Frame-wise thermal images—aligned using ThermoRAFT-AFP motion estimates—are fused into a continuous thermal mosaic. Frame spacing inherently reflects local deposition velocity, while feathered blending suppresses overlap artifacts and attenuates radiometric noise. Transparent dashed white lines indicate frame centers for spatial reference. For clarity, only a representative course segment is shown.

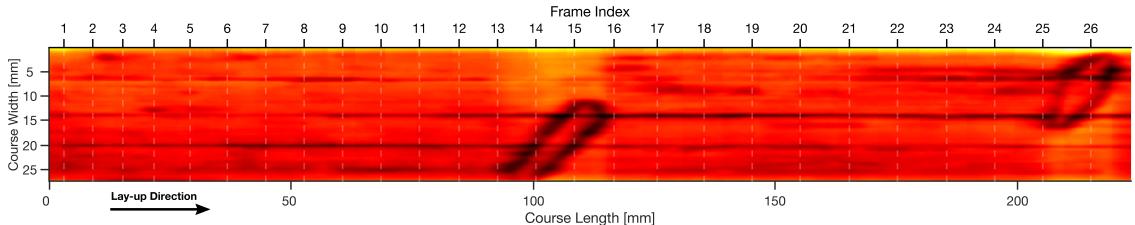


Figure 6.8 Course-level reconstruction for a uni-directional laminate containing both healthy and FOD-affected regions. Stitched frame centers are marked by transparent dashed white lines for reference. Only a representative segment of the full course is shown for visual clarity and space efficiency.

Global laminate reconstruction is achieved by stacking stitched courses. As shown in Figure 6.9, ply-wise synchronization is performed via cross-correlation of velocity profiles, compensating for inter-course shifts. In the illustrated example, a vertical lag of 12.3 mm is detected and corrected—ensuring seamless alignment across course boundaries. The resulting 2D lamina view enables full-ply inspection and supports inter-course defect interpretation.

To evaluate stitching fidelity, we compute the Structural Similarity Index Measure (SSIM) between the original frames and their corresponding regions in the course mosaic. An average SSIM of 0.8665 indicates strong structural preservation. Ply-

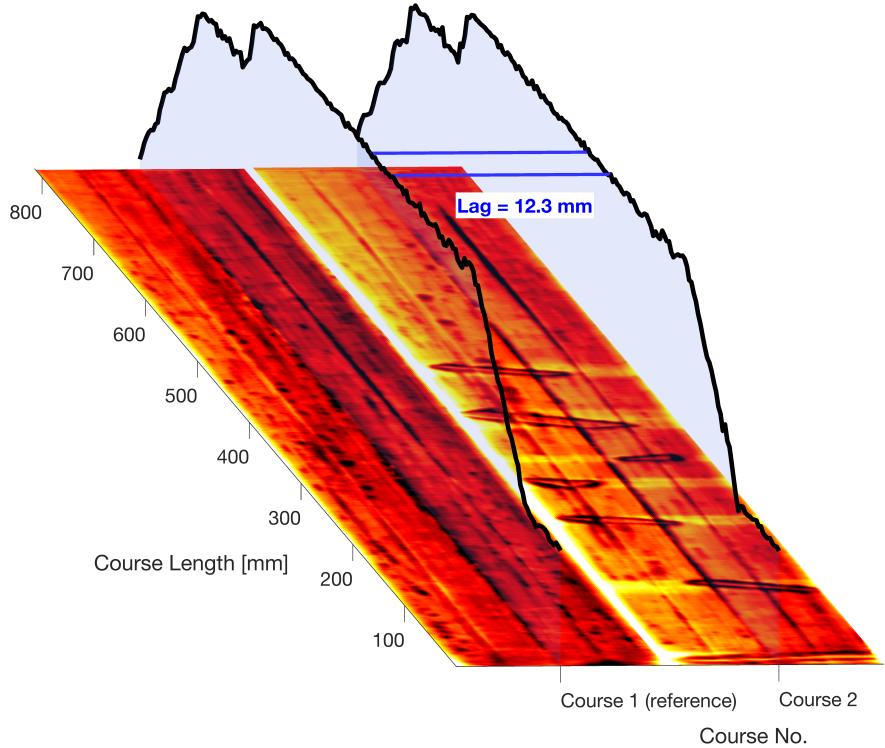


Figure 6.9 Ply-level motion synchronization visualization across two stitched AFP courses. Overlayed motion profiles (blue) highlight instantaneous velocity variations along the course length. Estimated inter-course lag is shown with a scaled marker in physical units (mm), derived from cross-correlation.

level registration accuracy is assessed by measuring vertical deviation at layer edges, yielding a mean absolute alignment error below 1.23 mm—achieved without any post-hoc correction.

Runtime profiling confirms real-time feasibility. Course-level stitching—including motion estimation, inter-frame alignment, lag correction, and blending—executes in 36.5 ms per frame, satisfying the 25 Hz (40 ms) throughput requirement. Ply-wise stacking operates asynchronously during  $\geq 5$  s machine idle intervals and introduces no delay to the real-time pipeline.

Unlike prior work such as (Denkena et al. (2022)), which focuses on per-frame motion, our method reconstructs entire laminates in real time. The output provides a continuous, operator-interpretable thermal map where localized anomalies can be understood in global context. This real-time, camera-only reconstruction supports cross-layer, inter-course, and intra-course defect analysis—marking a significant advance in in-situ inspection for automated composite manufacturing.

### 6.3.4 Reliability and Robustness Analysis

To evaluate the deployability of ThermoRAFT-APP under real-world AFP conditions, we assess its reliability and robustness across four dimensions: statistical agreement, variation across layers and defects, noise resilience, and tuning-free generalization.

**Agreement with Ground Truth:** A Bland–Altman analysis (Figure 6.10) is conducted over the full dataset to assess estimator reliability. The plotted subset (one-third for clarity) shows over 91.8% of frame-wise displacements fall within subpixel bounds ( $\pm 124 \mu\text{m}$ ), and 95% lie within  $\pm 1.96\sigma$  ( $\pm 211 \mu\text{m}$ ). The few outliers coincide with rapid tool transitions, suggesting that deviations are tied to transient thermographic noise rather than estimator error. Concordance Correlation Coefficient (CCC) reaches 99.64%, confirming strong agreement with the identity line. The 95% Total Deviation Index (TDI<sub>95</sub>) is 9  $\mu\text{m}$ , reflecting tight estimator bounds and high temporal repeatability.

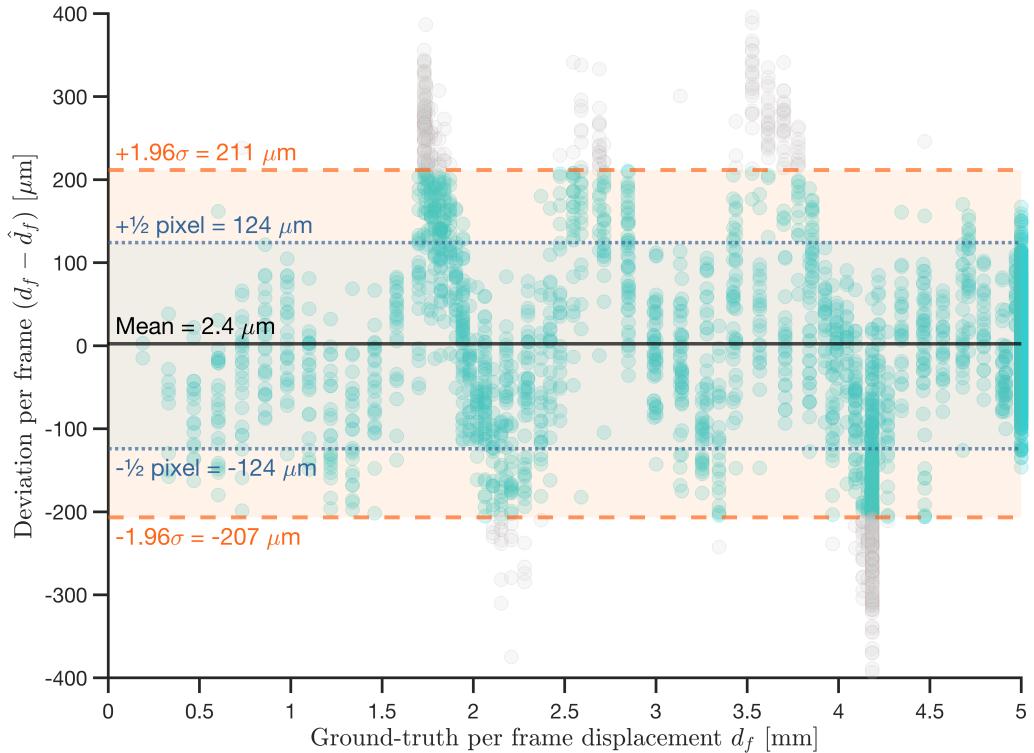


Figure 6.10 Bland–Altman analysis of frame-wise displacement estimates, showing deviation mean (black), 95% limits (orange), and  $\pm \frac{1}{2}$  pixel bounds (blue). Over 91% of estimates fall within the subpixel range; outliers coincide with velocity transitions.

**Performance Across Lay-up Depth and Texture:** As shown in Table 6.6, RMSE remains consistently low across all four ply levels, while MPE improves slightly with increasing lay-up depth. This suggests that estimator bias reduces

as thermal gradients stabilize. Texture-rich regions (e.g., defect zones) yield lower RMSE and more balanced MPE—indicating that the model effectively leverages spatial features without overfitting. The sign reversal in MPE between healthy (positive bias) and defective (slight underestimation) regions further illustrates stable estimator behavior across contrasting thermal inputs.

Table 6.6 Metric-wise comparison of ThermoRAFT-AFP performance across lay-up depth and feature richness. Shading reflects metric magnitude (lower is better). Arrows ( $\downarrow$ ) mark best values.  $\Delta H-D$  is the healthy-to-defective delta.

Metric	Layer No.				Mean	Lay-up Status		$\Delta H-D$
	1	2	3	4		Healthy	Defective	
RMSE [mm/s]	4.4	4.5	4.5	<b>4.3<math>\downarrow</math></b>	4.43	4.6	<b>4.2<math>\downarrow</math></b>	-0.4
MPE [%]	0.30	0.24	0.23	<b>0.17<math>\downarrow</math></b>	0.24	0.78	<b>-0.69<math>\downarrow</math></b>	1.47

**Noise Resilience Under Perturbations:** We simulate real-world thermographic noise via a Monte Carlo approach. Additive white Gaussian noise (AWGN), scaled to match FLIR A655sc’s NETD of 30 mK, is combined with a low-frequency columnar drift representing thermal shading. Perturbations are applied pre-normalization in the temperature domain to preserve realism.

As summarized in Table 6.7, ThermoRAFT-AFP maintains strong accuracy and temporal coherence across SNR levels down to 24.4 dB, with RMSE  $\approx$  4.68 mm/s, drift  $\approx$  0.44%, and  $R^2 > 99.4\%$ . Performance degrades gracefully at 14.4 dB (RMSE 7.05 mm/s, drift 1.18%,  $R^2$  86.9%), with functional limits observed near 4.3 dB. These results confirm robust estimator behavior under radiometric perturbations and bounded degradation under severe thermal noise.

**Tuning-Free Generalizability:** All results were achieved using a single, fixed configuration. Apart from a one-time regression calibration step, no re-training or scene-specific tuning was required. ThermoRAFT-AFP thus demonstrates strong cross-case generalization and high portability to varied AFP thermal setups—meeting a critical requirement for industrial applicability.

### 6.3.5 Benchmarking of Motion Estimation Methods

Reliable benchmarking for AFP thermal inspection is hindered by the absence of public datasets and limited reproducibility of prior work. To enable a fair and transparent comparison, we selected three representative motion estimation approaches:

Table 6.7 Metrology robustness of ThermoRAFT-AFP across noise levels ( $\sigma^2$ , SNR). Color shading qualitatively reflects performance drop per metric. The 14.4 dB row marks the final noise condition with acceptable estimation fidelity.

$\sigma^2$ (SNR)	RMSE	MAE	MPE	$R^2$	$\Delta_{\text{drift}}$
	[mm/s]	[mm/s]	[%]	[%]	[mm / %]
$10^{-3}$ (64.4 dB)	4.68	3.29	-0.67	99.46	3.36 / 0.44
$10^{-2}$ (54.4 dB)	4.68	3.29	-0.67	99.46	3.37 / 0.44
$10^{-1}$ (44.4 dB)	4.68	3.29	-0.67	99.46	3.37 / 0.44
$10^0$ (34.4 dB)	4.71	3.33	-0.68	99.45	3.39 / 0.45
$10^1$ (24.4 dB)	4.78	3.44	-0.67	99.43	3.28 / 0.43
$10^2$ (14.4 dB)	<b>7.05</b>	<b>5.66</b>	<b>2.21</b>	<b>86.92</b>	<b>8.96 / 1.18</b>
$10^3$ (4.3 dB)	395.34	248.93	165.28	-70.39	916.10 / 120.54

a dense classical method (Farneback (Farnebäck (2003))), a sparse keypoint tracker (SIFT with RANSAC (Lowe (2004))), and a domain-specific optical flow variant (CurvatureSSD), re-implemented from (Denkena et al. (2022)). All methods were evaluated on the same thermal AFP dataset under identical preprocessing, with parameters tuned via grid search. RAFT baseline results were previously analyzed in subsection 6.3.1 and are omitted here for conciseness.

Table 6.8 Comparison of motion estimation methods on thermal AFP dataset. Metrics span accuracy, runtime, and deployability. Bold marks best scores.

Method	Accuracy				Efficiency			Deployability		
	RMSE	MPE	$R^2$	$\Delta_{\text{drift}}$	FPS / Time / Budget	Real-Time	CCC / TDI <sub>95</sub>	Robustness	Generalizability	
	[mm/s]	[%]	[%]	[mm / %]	[Hz / ms / ms]		[% / $\mu\text{m}$ ]			
Farneback	60.28	-14.72	10.91	130.59 / 17.16	50 / 19.0 / 20.0	✓	59.7 / 169	Low	Low	
SIFT + RANSAC	7.66	3.18	98.53	18.82 / 2.48	25 / 9.4 / 40.0	✓✓	98.2 / 16	Moderate	Low	
CurvatureSSD†	7.38	-0.38	98.66	4.85 / 0.64	50 / 52.8 / 20.0	-	98.8 / 14	Moderate	Moderate	
<b>ThermoRAFT-AFP</b>	<b>4.83</b>	<b>0.17</b>	<b>99.33</b>	<b>0.67 / 0.09</b>	<b>25 / 33.0 / 40.0</b>	<b>✓</b>	<b>99.6 / 9</b>	<b>High</b>	<b>High</b>	

† CurvatureSSD re-implemented from algorithmic description due to unavailable source code or metrics.

As summarized in Table 6.8, ThermoRAFT-AFP surpasses all baselines across metrological accuracy, runtime efficiency, and field deployability. It achieves the lowest RMSE (4.83 mm/s), near-zero MPE, and minimal drift (0.09%), while maintaining  $R^2$  above 99.3%. In contrast, Farneback—despite operating at 50 FPS—exhibits extreme error and negative bias, rendering it unfit for AFP inspection tasks.

SIFT+RANSAC provides moderate accuracy with low latency (9.4 ms), but struggles under low-texture or high-noise inputs. CurvatureSSD performs well under

clean conditions but violates real-time constraints with 52.8 ms latency, and shows instability at reduced frame rates. Its evaluation excludes defects, cross-plies, and noise—raising concerns about deployment readiness.

By contrast, ThermoRAFT-AFP maintains 25 FPS real-time operation with 33.0 ms latency, subpixel accuracy, and robust performance across thermal, geometric, and structural variations. It also reports the highest CCC (99.64%) and lowest TDI<sub>95</sub> (9  $\mu$ m), evidencing consistent estimator agreement and generalization—all without scene-specific tuning.

These results affirm ThermoRAFT-AFP as the only solution satisfying the full spectrum of industrial requirements: accuracy, real-time capability, robustness to disturbance, and operational generalizability for thermal motion tracking in AFP monitoring.

#### 6.4 System-Level Qualitative Comparison

Table 6.9 presents a comprehensive system-level benchmark of recent AFP inspection studies, evaluated across 13 core functional capabilities—ranging from basic in-situ sensing to advanced motion estimation and laminate-scale reconstruction. This capability matrix is designed to assess the functional maturity of inspection solutions, highlighting the degree of integration and operational readiness. It serves not only as a landscape overview but also as a critical positioning tool for the proposed work within both thermal and laser-based research streams.

A first inspection reveals that most prior studies remain confined to narrow functional slices. For example, (Meister & Wermes (2023)) and (Sacco et al. (2020)) focus solely on defect detection and classification. Others, such as (Ghamisi et al. (2023)), introduce tow identification but stop short of segmentation, localization, or any form of thermal motion analysis. Even among the most comprehensive works, no prior study satisfies more than 6 of the 13 capability categories—underscoring the fragmented state of the field.

Notably, only a few studies such as (Schmidt et al. (2019)) and (Zhang et al. (2022)) demonstrate real-time operation, yet both omit critical modules such as motion estimation and reconstruction. Similarly, works like (Tang et al. (2022)) and (Sacco et al. (2020)) achieve segmentation and classification but lack global consistency or

Table 6.9 System-level functional coverage comparison of AFP inspection systems across recent studies. Capabilities span defect analytics, thermal motion tracking, and real-time deployment.

Reference	Sensor	Methodology	Database	In-Situ Monitoring	Tow Identification	Defect Detection	Defect Classification	Defect Segmentation	Defect Localization	Quality Evaluation	Motion Estimation	Lay-up Reconstruction	Real-Time Operation
Meister & Wermes (2023)	laser	data-driven	469	-	-	✓	✓	-	-	-	-	-	-
Zhang, Wang, Liu, Guo & Ji (2022)	laser	data-driven	3,000	-	-	✓	✓	-	✓	✓	-	-	✓
Tang et al. (2022)	laser	hybrid	43	✓	-	✓	✓	✓	✓	-	-	-	✓
Sacco et al. (2020)	laser	data-driven	800	-	-	✓	✓	✓	✓	-	-	-	-
Ghamisi et al. (2023)	laser	data-driven	44	-	✓	✓	-	-	✓	-	-	-	-
Schmidt et al. (2019)	thermal	data-driven	12,000	✓	-	✓	✓	-	-	-	-	-	✓
Juarez & Gregory (2021)	thermal	model-based	16	✓	-	-	-	-	-	✓	-	✓	✓
Denkena et al. (2022)	thermal	model-based	~250	✓	-	-	-	-	-	✓	✓	-	-
<b>This work</b>	thermal	hybrid	13,300 <sup>†</sup>	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

<sup>†</sup> 13,300 frames used for motion estimation and reconstruction; 5,000 subset used for defect-level tasks.

cross-layer inspection, as they operate on frame-wise analysis alone.

The motion estimation study by (Denkena et al. (2022)) provides an important exception. It introduces thermal optical flow for lay-up tracking but does so under restrictive conditions: healthy layers only, no segmentation, and no real-time or reconstruction support. Consequently, it remains limited in scope and offers no direct means for defect-level inspection or decision support.

In contrast, the proposed system—highlighted in the shaded row—stands out as the only approach offering complete functional integration. It covers all 13 capabilities across both data-driven and model-based domains, from dynamic tow identification to defect localization, thermal motion estimation, and real-time laminate reconstruction. Notably, these are achieved without machine-side feedback, and at operational frame rates suitable for on-line deployment.

This comprehensive functionality is enabled by a dual-framework architecture. The first stream focuses on localized, frame-wise analysis using hybrid learning algorithms for defect detection, classification, segmentation, and tow-aware localization. The second stream reconstructs temporally consistent lay-up mosaics using dense thermal motion estimation and synchronization—facilitating global quality evaluation. Combined, they produce a unified, learning-centric platform that transitions from isolated thermal frames to interpretable, ply-level visualizations for operator assistance.

Importantly, these results are achieved using an internal thermal dataset of 13,300 frames—substantially larger than those used in comparable studies—and validated across both healthy and defective lay-ups. The subset of 5,000 frames used for defect-level tasks ensures sufficient diversity and realism, while preserving scalability and repeatability.

In summary, this work delivers the first fully operational, real-time, thermography-based AFP inspection system with end-to-end functional coverage. It not only closes long-standing gaps in segmentation, tracking, and reconstruction, but also establishes a practical framework for future industrial deployment—marking a step change in the state of in-situ composite process monitoring.

## 7. INTEGRATED SYSTEM OPERATION: PRACTICAL CASE

### STUDIES DEMONSTRATION

Having validated both frameworks quantitatively, we now demonstrate their integration under realistic conditions. This chapter presents the real-time operation of the proposed dual-framework inspection system through two representative AFP lay-up scenarios. Engineered for industrial relevance and deployability, the system unifies frame-wise defect analysis with global thermal reconstruction into a cohesive, visually rich, and operator-assistive pipeline. The two frameworks operate asynchronously: motion estimation and lay-up reconstruction run at 25 frames per second (FPS), enabling course- and laminate-level visualization, while defect analysis and quality assessment operate at 5 FPS for focused, non-overlapping frame-wise inspection. Despite this disparity, outputs are fused within a unified interface that evolves incrementally as new frames are acquired, supporting both real-time monitoring and historical review.

The envisioned GUI overlays local defect outputs onto the global thermal view in real time. Operators can inspect reconstructed laminates, trace anomalies to individual frames or courses, and receive actionable recommendations based on metrics such as the DAP ratio. To demonstrate this layered functionality, two case studies are presented: the first involves a unidirectional laminate with multiple FODs; the second addresses a cross-ply lay-up featuring various defects such as gaps, overlaps, and an organic tow splice. Together, these scenarios validate the system's capability for interpretable, layered, and machine-independent thermographic AFP inspection in realistic production settings.

## 7.1 Case Study A: Unidirectional Lay-up with Multiple FODs

The first scenario demonstrates system behavior under a unidirectional lay-up contaminated with multiple FODs, such as tow scrap and airborne particulates. The laminate configuration follows a typical  $[0^\circ, 0^\circ, 0^\circ, 0^\circ]$  stacking sequence, representative of aerospace fuselage floors and stiffener webs. Contamination was manually introduced in the first two courses of the second ply (C1–C2 in L2) to emulate realistic production faults in high-throughput AFP settings.

The reconstructed laminate, visualized in Figure 7.1, reveals several structurally relevant patterns not captured in conventional frame-wise views. Most notably, elongated heat trails perpendicular to the tow direction appear at nearly every FOD site in ply 2, indicating potential partial disbondment or tack failure. These patterns extend vertically into ply 3 and attenuate into ply 4, highlighting a volumetric thermal influence shaped by defect geometry, material properties, and ply thickness.

The presence of cross-tow thermal spread also suggests the feasibility of augmenting the current DAP metric with directional propagation cues, enabling future variants sensitive to such tack failure modes. This depth-aware thermal view enables operators to assess not only surface anomalies but also cross-layer propagation trends—shifting inspection logic from isolated frame-level detection to full-ply thermal diagnostics grounded in global coherence.

An example frame from the contaminated region is processed by the frame-wise defect analysis pipeline, as shown in Figure 7.2. After preprocessing, the input is flagged as defective by the SVM-based binary classifier. The defect is first classified using the DCNN-based classifier into the FB defect class, and the defect region is then segmented using an active contours algorithm and overlaid on the thermal frame with a bounding box for spatial localization.

The system computes DAP values at both course and tow levels, with the most severe case—tow 4—reaching a value of 70.56%. Since this value exceeds the expert-defined 40% severity threshold, the outcome triggers an alert from the predefined AFP-DSS knowledge base. This alert–recommendation pair is derived from a combination of qualitative indicators (defect class, ply number, tow index) and quantitative metrics (DAP values), and recommends stopping the process and removing the foreign object.

This case illustrates how the integrated system synchronizes global reconstruction with local detection to deliver interpretable and actionable defect intelligence. The

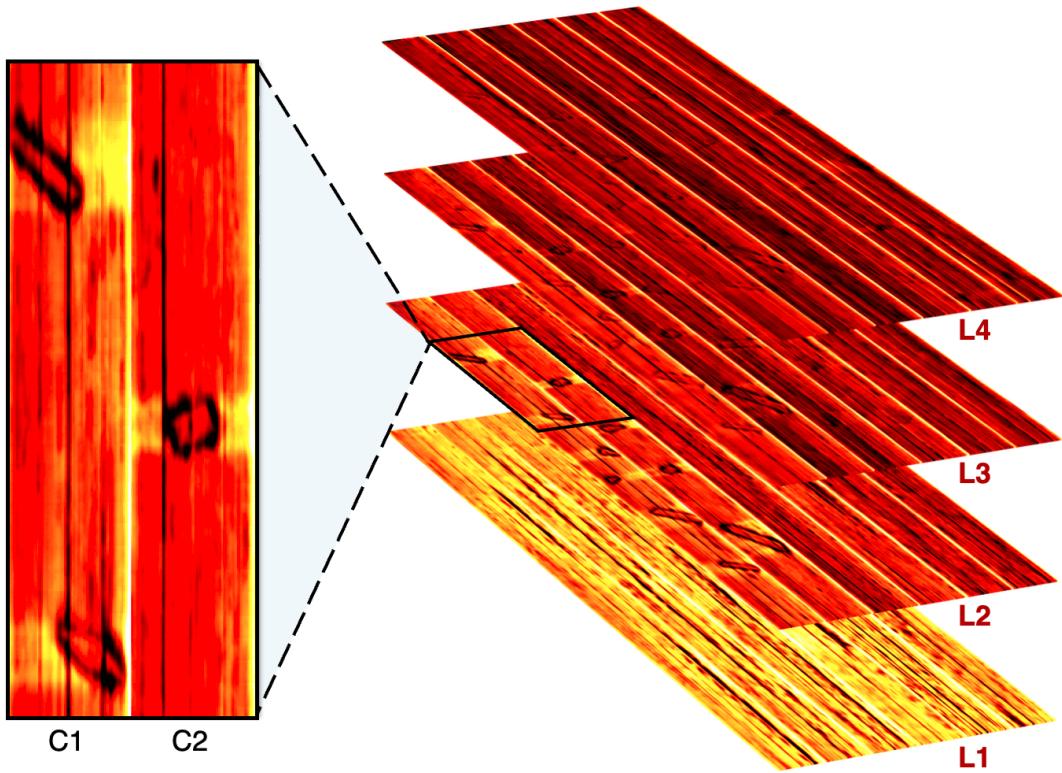


Figure 7.1 Case A: Unidirectional [0°] laminate with tow scrap FODs. Right: exploded-view reconstruction showing vertical thermal propagation. Left: zoom-in on ply 2 highlighting lateral thermal separation zones.

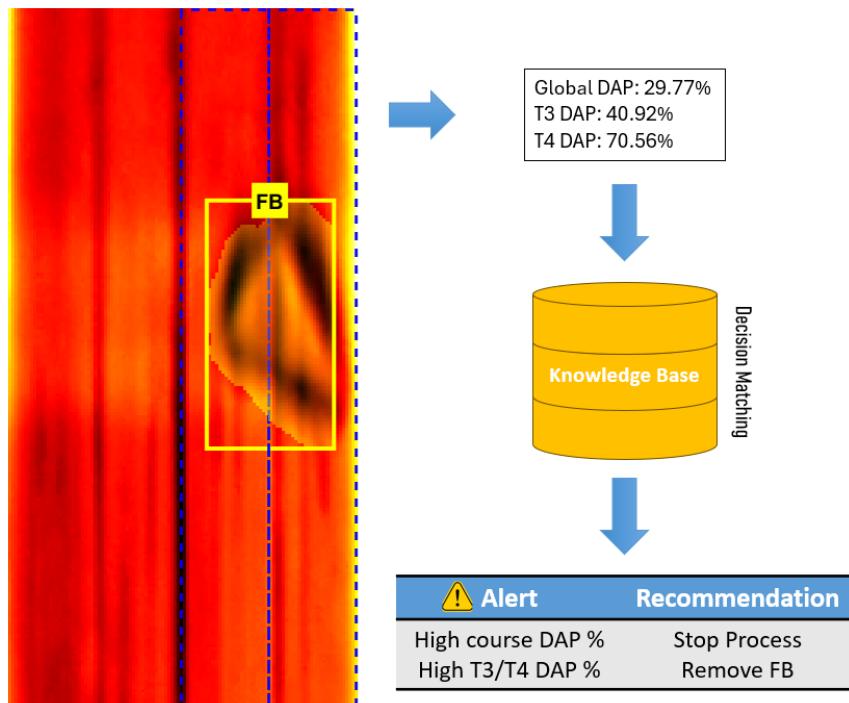


Figure 7.2 Frame-wise inference on a defective instance from Case A. Thermal anomaly is segmented and localized with high tow-wise DAP. The decision system recommends process stoppage and defect removal.

observed alignment between ply-level thermal trails and frame-level segmentation reinforces the benefits of multi-scale fusion for reliable, operator-assistive AFP quality assurance.

## 7.2 Case Study B: Cross-Ply Lay-up with Mixed Defect Types

The second scenario evaluates the system's robustness under structurally complex and thermally diverse conditions by examining a balanced cross-ply laminate with mixed defect types. The lay-up sequence follows a  $[0^\circ, 90^\circ, 0^\circ, 90^\circ]$  configuration, representative of quasi-isotropic aerospace components such as wing skins and pressure bulkheads. This arrangement introduces directional variation in fiber orientation, thermal texture, and potential defect manifestations—presenting a more demanding inspection context.

The reconstructed laminate, shown in Figure 7.3, reveals multiple defect signatures across plies, visualized through an exploded 3D representation. Notably, three anomalies are spatially resolved: (a) an embedded tow splice in ply 3, (b) a localized delamination trail, and (c) cut-induced gaps in plies 2 and 3, occasionally accompanied by overlaps. Despite orthogonal ply stacking and motion complexity, cumulative drift remains under 1% of total course length, validating the stability of the real-time motion estimation and stitching framework. This ensures both machine-independent automation and reliable operator-guided traceability.

A single example frame is analyzed via the frame-wise inspection framework, as shown in Figure 7.4. The frame is processed through the universal pipeline, with the classifier identifying the anomaly as a longitudinal gap centered on tow 3. The segmented region is localized and overlaid using active contours, while DAP scores are computed to assess severity. Tow 3 exhibits a high DAP of 45.08%, exceeding the expert-defined threshold for process intervention.

The integrated decision layer then matches the outcome to a predefined entry in the AFP-DSS knowledge base. Based on both qualitative attributes (defect class, affected tow, ply index) and quantitative scores, the system recommends pausing the process and adjusting or replacing tow 3 before resuming. Notably, for narrower gaps that fall within acceptable limits, the system can defer to autoclave-induced autocorrection, thereby avoiding unnecessary interruptions.

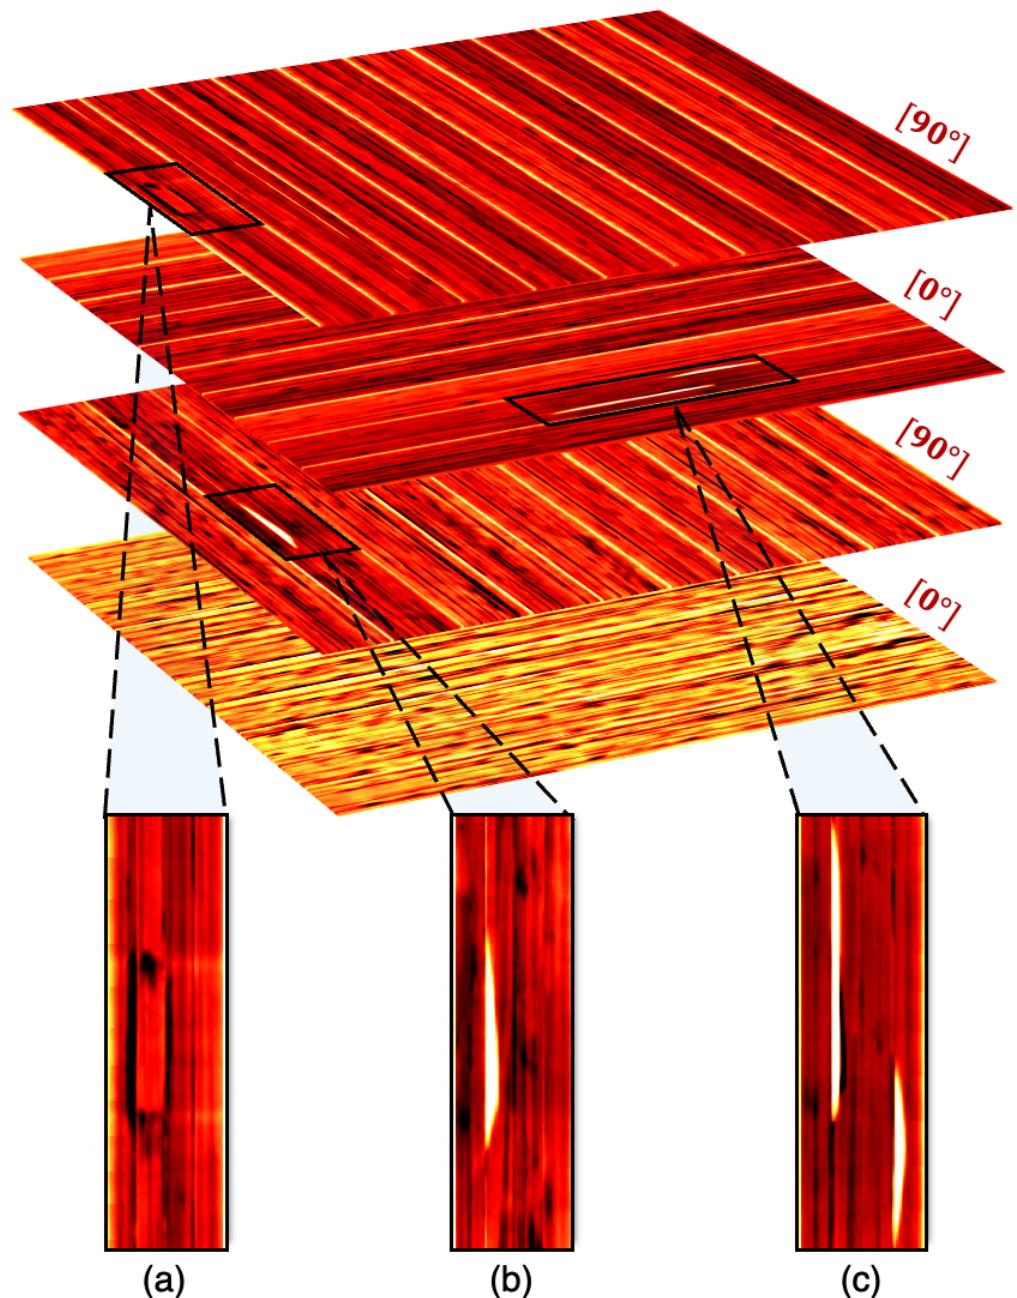


Figure 7.3 Case B: Cross-ply  $[0^\circ, 90^\circ, 0^\circ, 90^\circ]$  laminate with mixed defect types. Exploded 3D view reveals (a) tow splice, (b) delamination trail, and (c) cut-induced gap and accompanying overlap, demonstrating system accuracy under directional variation.

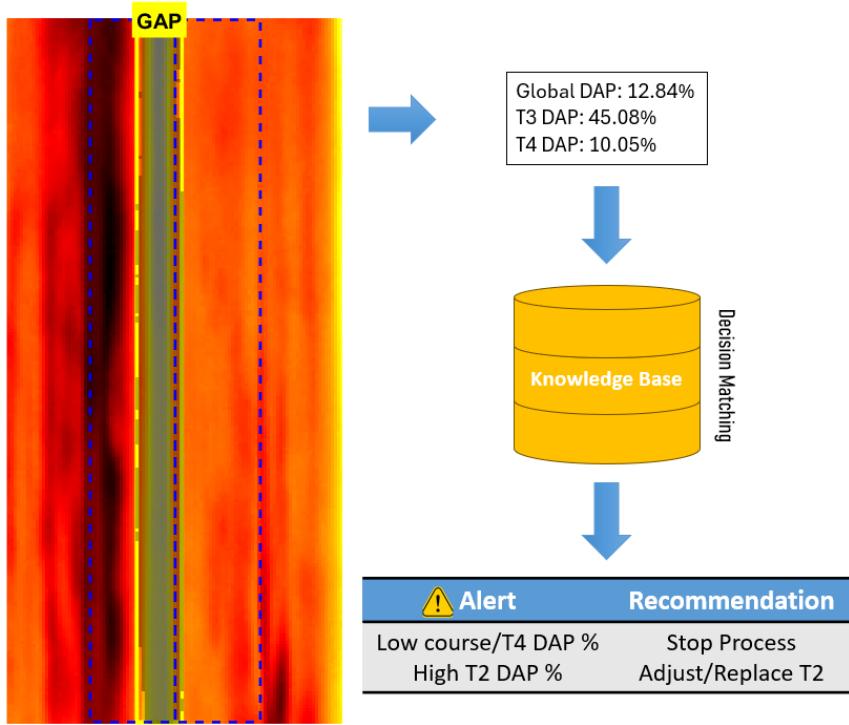


Figure 7.4 Frame-wise inference on a defective instance from Case B. Localized tow gap triggers DAP-based alert and recommends corrective action via the knowledge base.

This case yields two key observations. First, defects that preserve material continuity—such as overlaps and splices—tend to exhibit stronger cross-ply thermal propagation compared to void-type anomalies like gaps. Second, the combined use of global reconstruction and local inspection enables the system to differentiate transient surface noise from thermally coherent volumetric defects. Together, these capabilities reinforce the value of interpretable, learning-centric, and thermally grounded AFP inspection across diverse lay-up configurations.

## 8. CONCLUSIONS AND FUTURE DIRECTIONS

This dissertation presented a learning-centric, end-to-end hybrid inspection system for real-time quality assessment in Automated Fiber Placement (AFP), significantly advancing the state-of-the-art in thermal vision-based process monitoring. The proposed dual-framework architecture—uniting a hybrid frame-wise defect analysis and quality assessment module with a machine-independent, motion-aware global reconstruction pipeline—delivers both localized defect intelligence and laminate-scale traceability. This synergistic combination addresses longstanding limitations in manual and automated inspection workflows by enabling timely, interpretable, and scalable analysis of AFP lay-ups through a purely thermographic solution.

The system eliminates operator dependency by automating key inspection tasks, including dynamic tow boundary identification, defect detection, classification, segmentation, and lay-up quality evaluation. Central to this capability are several algorithmic innovations: a hierarchical Gabor-filter-based classification cascade, a texture-guided active contour segmentation algorithm, and a novel DAP metric quantifying defect severity at both tow and course levels. These modules operate within a unified, setup-independent framework and were rigorously validated across diverse thermal datasets representative of real-world AFP production conditions.

Complementing this local inspection framework, ThermoRAFT-AFP—the second core module—delivers dense thermal motion estimation through a tailored adaptation of RAFT optical flow. AFP-specific augmentations such as predictive warm-start initialization, drift correction, and temporal stabilization enable subpixel-accurate flow estimation in real time. Operating at 25 fps, the system reconstructs course-wise and ply-wise laminate visualizations with high spatial fidelity and cumulative drift below 1%. Evaluations over 13,000+ thermal frames confirm the system’s robustness and generalizability, achieving velocity RMSE below 5 mm/s, mean percentage error of 0.17%, and F1-scores exceeding 96% for defect classification—all without requiring any machine-side data or encoder feedback.

A core outcome of this work is the seamless integration of both frameworks within

an operator-assistive, vision-only inspection interface. Frame-wise analytics and global reconstructions are fused into a dynamic graphical environment, empowering users to monitor, trace, and respond to defects across spatial and temporal scales. The resulting system functions as a comprehensive AFP-DSS—an AFP Decision Support System—capable of translating complex thermal patterns into actionable quality intelligence. Through its embedded alert and recommendation engine, the system enables real-time diagnostics, minimizes machine downtime, and promotes informed decision-making across the entire lay-up cycle.

The proposed solution satisfies all core industrial criteria: (i) real-time operability on CPU-based hardware, (ii) interpretability through visual overlays and calibrated metrics, (iii) robustness to lay-up variations and defect diversity, and (iv) scalability to large laminates via motion-aware stitching. Its modular structure supports both live process monitoring and retrospective quality audit modes, enabling broad applicability and easy integration with existing AFP workflows.

Nevertheless, several limitations remain. The system has thus far been validated primarily under fixed-stiffness production regimes and may not generalize directly to variable-stiffness or curvilinear lay-ups. While most learning-based components—such as classification and segmentation—can be extended to these scenarios with appropriate retraining, other elements, including the motion estimation pipeline and preprocessing routines, may require substantial re-engineering to accommodate non-uniform fiber orientations and complex surface geometries. Additionally, relying solely on thermal imaging, though beneficial for deployment simplicity, may constrain classification accuracy in ambiguous or low-contrast cases where multi-modal sensing could provide enhanced discrimination. Lastly, while the DAP metric offers a robust thermal-based estimate of defect severity, further studies are needed to correlate its values with the mechanical performance of the final composite part.

Overall, this dissertation contributes a foundational step toward intelligent, scalable, and deployable AFP inspection systems. It bridges the gap between thermal vision, hybrid AI, and industrial inspection by unifying high-frequency analytics with human-interpretable, laminate-level visualization—paving the way for resilient, autonomous composite manufacturing workflows.

## Future Work

Building upon the success of the proposed system, several research avenues hold strong potential for extending its capabilities. In the short term, the defect analysis pipeline can be enriched to support a wider range of defect types—including wrinkles, bridging, tow twists, and surface contamination—through retraining on expanded and diverse labeled datasets. Integrating inter-ply defect tracking into the reconstruction process will enable temporal propagation analysis, providing deeper insight into multi-layer defect evolution. Moreover, calibration of the DAP metric against mechanical testing would allow the system to estimate structural impact, enhancing its utility in predictive diagnostics. Operator-facing improvements—such as a feedback-enabled GUI for annotation, review, and traceability—will support both user trust and data-driven iterative learning.

On a broader horizon, the framework can be extended to handle non-planar, curvilinear, and variable-stiffness lay-ups by incorporating geometry-aware motion estimation, curvature-adaptive segmentation, and non-uniform stitching logic. Depth sensing, stereo vision, or shape-from-motion techniques may be integrated to enable 3D surface reconstruction, improving detection of out-of-plane anomalies such as tow lift, bridging, and buckling. Robustness under challenging conditions can be further enhanced by exploring thermal–visual fusion models, enabling more reliable defect classification in noisy or ambiguous regions. A unified, lightweight multi-task learning model—capable of simultaneously performing motion estimation, classification, and segmentation—could streamline real-time inference, particularly for edge-based or embedded deployments.

In the long term, the vision-centric design of the proposed system makes it a strong candidate for integration within smart, connected AFP environments. Closing the loop with AFP machine controllers would enable autonomous repair, tow re-application, and adaptive process adjustments based on real-time defect detection—transforming the system from a passive observer into an active controller. Scaling the framework into a cloud-enabled, networked platform could facilitate dataset-driven optimization, cross-machine performance benchmarking, and holistic laminate quality traceability across entire production lines—advancing the frontier of intelligent, adaptive, and autonomous composite manufacturing.

## BIBLIOGRAPHY

Abouhamzeh, M., Nardi, D., Leonard, R., & Sinke, J. (2018). Effect of prepreg gaps and overlaps on mechanical properties of fibre metal laminates. *Composites Part A: Applied Science and Manufacturing*, 114, 258–268.

Alfarano, A., Maiano, L., Papa, L., & Amerini, I. (2024a). Estimating optical flow: A comprehensive review of the state of the art. *Computer Vision and Image Understanding*, 249, 104160.

Alfarano, A., Maiano, L., Papa, L., & Amerini, I. (2024b). Estimating optical flow: A comprehensive review of the state of the art. *Computer Vision and Image Understanding*, 104160.

Anay, R., Miller, D., Tessema, A., Wehbe, R., Ziehl, P., Tatting, B., Gurdal, Z., Harik, R., & Kidane, A. (2022). An experimental investigation concerning the effect of afp defects on progressive damage in cfrp coupons. *Composite Structures*, 279, 114725.

August, Z., Ostrander, G., Michasiow, J., & Hauber, D. (2014). Recent developments in automated fiber placement of thermoplastic composites. *SAMPE J*, 50(2), 30–37.

Bindu, C. H. & Prasad, K. S. (2012). An efficient medical image segmentation using conventional otsu method. *International Journal of Advanced Science and Technology*, 38(1), 67–74.

Blom, A. W., Lopes, C. S., Kromwijk, P. J., Gurdal, Z., & Camanho, P. (2009). A theoretical model to study the influence of tow-drop areas on the stiffness and strength of variable-stiffness laminates. *Journal of Composite Materials*, 43(5), 403–425.

Bradski, G. (2000). The OpenCV Library. *Dr. Dobb's Journal of Software Tools*.

Brasington, A., Sacco, C., Halbritter, J., Wehbe, R., & Harik, R. (2021). Automated fiber placement: A review of history, current technologies, and future paths forward. *Composites Part C: Open Access*, 6, 100182.

Butler, D. J., Wulff, J., Stanley, G. B., & Black, M. J. (2012). A naturalistic open source movie for optical flow evaluation. In *Computer Vision–ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7–13, 2012, Proceedings, Part VI 12*, (pp. 611–625). Springer.

Canny, J. (1986). A computational approach to edge detection. *Pattern Analysis and Machine Intelligence, IEEE Transactions on, PAMI-8*, 679 – 698.

Carosella, S., Hügle, S., Helber, F., & Middendorf, P. (2024). A short review on recent advances in automated fiber placement and filament winding technologies. *Composites Part B: Engineering*, 287, 111843.

Cemenska, J., Rudberg, T., & Henscheid, M. (2015). Automated in-process inspection system for afp machines. *SAE International Journal of Aerospace*, 8(2), 303–309.

Cervantes, J., Garcia-Lamont, F., Rodríguez-Mazahua, L., & Lopez, A. (2020). A comprehensive survey on support vector machine classification: Applications, challenges and trends. *Neurocomputing*, 408, 189–215.

Chan, T. F. & Vese, L. A. (2001). Active contours without edges. *IEEE Transactions on image processing*, 10(2), 266–277.

Denkena, B., Schmidt, C., Timmermann, M., & Friedel, A. (2022). An optical-flow-based monitoring method for measuring translational motion in infrared-thermographic images of afp processes. *Production Engineering*, 16(4), 569–578.

Denkena, B., Schmidt, C., Völtzer, K., & Hocke, T. (2016). Thermographic online monitoring system for automated fiber placement processes. *Composites Part B: Engineering*, 97, 239–243.

Denkena, B., Schmidt, C., & Weber, P. (2016). Automated fiber placement head for manufacturing of innovative aerospace stiffening structures. *Procedia Manufacturing*, 6, 96–104.

Dosovitskiy, A., Fischer, P., Ilg, E., Hausser, P., Hazirbas, C., Golkov, V., Van Der Smagt, P., Cremers, D., & Brox, T. (2015). Flownet: Learning optical flow with convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, (pp. 2758–2766).

Drinkwater, B. W. & Wilcox, P. D. (2006). Ultrasonic arrays for non-destructive evaluation: A review. *NDT & e International*, 39(7), 525–541.

Farnebäck, G. (2003). Two-frame motion estimation based on polynomial expansion. In *Image Analysis: 13th Scandinavian Conference, SCIA 2003 Halmstad, Sweden, June 29–July 2, 2003 Proceedings* 13, (pp. 363–370). Springer.

Geiger, A., Lenz, P., Stiller, C., & Urtasun, R. (2013). Vision meets robotics: The kitti dataset. *The international journal of robotics research*, 32(11), 1231–1237.

Ghamisi, A., Charter, T., Ji, L., Rivard, M., Lund, G., & Najjaran, H. (2023). Anomaly detection in automated fibre placement: Learning with data limitations. *arXiv preprint arXiv:2307.07893*.

Gonzalez, R. C. (2009). *Digital image processing*. Pearson education india.

Gregory, E. D. & Juarez, P. D. (2018). In-situ thermography of automated fiber placement parts. *AIP Conference Proceedings*, 1949(1), 060005.

Hanbay, K., Talu, M. F., & Özgüven, Ö. F. (2016). Fabric defect detection systems and methods—a systematic literature review. *Optik*, 127(24), 11960–11973.

Hartley, R. & Zisserman, A. (2003). *Multiple view geometry in computer vision*. Cambridge university press.

Heikkila, J. & Silvén, O. (1997). A four-step camera calibration procedure with implicit image correction. In *Proceedings of IEEE computer society conference on computer vision and pattern recognition*, (pp. 1106–1112). IEEE.

Hilsenstein, V. (2005). Surface reconstruction of water waves using thermographic stereo imaging. In *Image and Vision Computing New Zealand*, volume 2, (pp. 102–107). Citeseer.

Hocke, T. (2020). *Classification and Investigation of Thermographically Recorded Manufacturing Defects in the Automated Fiber Placement Process*. Ph.d. dissertation, Gottfried Wilhelm Leibniz Universität Hannover. In German.

Ilg, E., Mayer, N., Saikia, T., Keuper, M., Dosovitskiy, A., & Brox, T. (2017). Flownet 2.0: Evolution of optical flow estimation with deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Jähne, B., Scharr, H., Körkel, S., Jähne, B., Haußecker, H., & Geißler, P. (1999). Principles of filter design. *Handbook of computer vision and applications*, 2, 125–151.

Javed, A., Mirza, A. U., et al. (2013). Comparative analysis of different fabric defects detection techniques. *International journal of image, Graphics and Signal Processing*, 5(1), 40.

Juarez, P., Cramer, E., & Seebo, J. (2016). Advances in in situ inspection of automated fiber placement systems. (pp. 986109).

Juarez, P. & Gregory, E. (2019). In situ thermal inspection of automated fiber placement manufacturing. volume 2102, (pp. 120005).

Juarez, P. & Gregory, E. (2021). In situ thermal inspection of automated fiber placement for manufacturing induced defects. *Composites Part B: Engineering*, 220, 109002.

Juarez, P. D., Gregory, E. D., & Cramer, K. E. (2018). In situ thermal inspection of automated fiber placement operations for tow and ply defect detection. Technical report.

Kalman, R. E. (1960). A new approach to linear filtering and prediction problems.

Koptelov, A., El Said, B., & Tretiak, I. (2025). Enhancing afp manufacturing with ai: Defects forecasting and classification. *Composites Part B: Engineering*, 304, 112655.

Kumar, A. (2008). Computer-vision-based fabric defect detection: A survey. *IEEE Transactions on Industrial Electronics*, 55(1), 348–363.

Lowe, D. G. (2004). Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60, 91–110.

Mahajan, P., Kolhe, S., & Patil, P. (2009). A review of automatic fabric defect detection techniques. *Advances in Computational Research*, 1(2), 18–29.

Martins Briedis, K., Gross, M., & Schroers, C. (2025). Efficient correlation volume sampling for ultra-high-resolution optical flow estimation. *arXiv e-prints*, arXiv–2505.

Meister, S., Möller, N., Stüve, J., & Groves, R. M. (2021). Synthetic image data augmentation for fibre layup inspection processes: Techniques to enhance the data set. *Journal of Intelligent Manufacturing*, 32, 1767–1789.

Meister, S. & Wermes, M. (2023). Performance evaluation of cnn and r-cnn based line by line analysis algorithms for fibre placement defect classification. *Production Engineering*, 17(3-4), 391–406.

Meister, S., Wermes, M., Stüve, J., & Groves, R. (2021). Review of image segmentation techniques for layup defect detection in the automated fiber placement process: A comprehensive study to improve afp inspection. *Journal of Intelligent Manufacturing*, 32.

Meister, S., Wermes, M. A. M., Stüve, J., & Groves, R. M. (2020). Algorithm assessment for layup defect segmentation from laser line scan sensor based image data. In *Sensors and smart structures technologies for civil, mechanical, and aerospace systems 2020*, volume 11379, (pp. 139–153). SPIE.

Ng, H., Du, R., et al. (2005). Acquisition of 3d surface temperature distribution of a car body. In *2005 IEEE International Conference on Information Acquisition*, (pp. 5–pp). IEEE.

Oromiehie, E., Prusty, B. G., Compston, P., & Rajan, G. (2019). Automated fibre placement based composite structures: Review on the defects, impacts and inspections techniques. *Composite Structures*, 224, 110987.

Oromiehie, E., Prusty, B. G., Rajan, G., & Compston, P. (2016). Optical fiber bragg grating sensors for process monitoring in advanced composites. In *2016 IEEE*

*Sensors Applications Symposium (SAS)*, (pp. 1–5). IEEE.

Otsu, N. (1979). A threshold selection method from gray-level histograms. *IEEE transactions on systems, man, and cybernetics*, 9(1), 62–66.

Pantofi, S., Kassapoglou, C., & Peeters, D. (2025). Predicting gaps and overlaps in automated fiber placement composites by measuring sources of manufacturing process variations. *Composites Part B: Engineering*, 291, 111891.

Parmar, H., Khan, T., Tucci, F., Umer, R., & Carbone, P. (2022). Advanced robotics and additive manufacturing of composites: towards a new era in industry 4.0. *Materials and manufacturing processes*, 37(5), 483–517.

Passing, H. & Bablok, W. (1983). A new biometrical procedure for testing the equality of measurements from two different analytical methods. application of linear regression procedures for method comparison studies in clinical chemistry, part i.

Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., & Lerer, A. (2017). Automatic differentiation in pytorch.

Rudberg, T., Cemenska, J., & Sherrard, E. (2019). A process for delivering extreme afp head reliability. *SAE International Journal of Advances and Current Practices in Mobility*, 1(2), 333–342.

Rudberg, T., Nielson, J., Henscheid, M., & Cemenska, J. (2014). Improving afp cell performance. *SAE International Journal of Aerospace*, 7(2), 317.

Sacco, C., Baz Radwan, A., Anderson, A., Harik, R., & Gregory, E. (2020). Machine learning in composites manufacturing: A case study of automated fiber placement inspection. *Composite Structures*, 250, 112514.

Sacco, C., Radwan, A. B., Beatty, T., & Harik, R. (2019). Machine learning based afp inspection: A tool for characterization and integration. *SAMPE 2019-Charlotte, NC, May 2019*.

Sacco, C., Radwan, A. B., Harik, R., & Van Tooren, M. (2018). Automated fiber placement defects: Automated inspection and characterization. In *SAMPE 2018 Conference and Exhibition*, number NF1676L-29116.

Savitzky, A. & Golay, M. J. (1964). Smoothing and differentiation of data by simplified least squares procedures. *Analytical chemistry*, 36(8), 1627–1639.

Schmidt, C., Denkena, B., Völtzer, K., & Hocke, T. (2017). Thermal image-based monitoring for the automated fiber placement process. *Procedia CIRP*, 62, 27–32. 10th CIRP Conference on Intelligent Computation in Manufacturing Engineering - CIRP ICME '16. [Edited by: Roberto Teti, Manager Editor: Doriana M. D'Addona].

Schmidt, C., Hocke, T., & Denkena, B. (2019). Deep learning-based classification of production defects in automated-fiber-placement processes. *Production Engineering*, 13, 501–509.

Shadmehri, F., Ioachim, O., Pahud, O., Brunel, J., Landry, A., Hoa, V., & Hojjati, M. (2015). Laser-vision inspection system for automated fiber placement (afp) process. In *20th International conference on composite materials Copenhagen*.

Spearman, C. (1961). The proof and measurement of association between two things.

Sreeshan, K., Dinesh, R., & Renji, K. (2020). Nondestructive inspection of aerospace composite laminate using thermal image processing. *SN Applied Sciences*, 2(11), 1830.

Sun, D., Yang, X., Liu, M.-Y., & Kautz, J. (2018). Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Sun, S., Han, Z., Fu, H., Jin, H., Dhupia, J. S., & Wang, Y. (2020). Defect characteristics and online detection techniques during manufacturing of frps using automated fiber placement: A review. *Polymers*, 12(6), 1337.

Tang, Y., Wang, Q., Cheng, L., Li, J., & Ke, Y. (2022). An in-process inspection method integrating deep learning and classical algorithm for automated fiber placement. *Composite Structures*, 300, 116051.

Teed, Z. & Deng, J. (2020). Raft: Recurrent all-pairs field transforms for optical flow. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, (pp. 402–419). Springer.

Turner, M. R. (1986). Texture discrimination by gabor functions. *Biological cybernetics*, 55(2), 71–82.

Van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., Boulogne, F., Warner, J. D., Yager, N., Gouillart, E., & Yu, T. (2014). scikit-image: image processing in python. *PeerJ*, 2, e453.

Wang, Y., Wang, W., Li, Y., Guo, J., Xu, Y., Ma, J., Ling, Y., Fu, Y., & Jia, Y. (2024). Research on traditional and deep learning strategies based on optical flow estimation - a review. *Journal of King Saud University - Computer and Information Sciences*, 36(4), 102029.

Yadav, N. & Schledjewski, R. (2023). Review of in-process defect monitoring for automated tape laying. *Composites Part A: Applied Science and Manufacturing*, 173, 107654.

Yipeng, T., Wang, Q., Wang, H., Li, J., & Ke, Y. (2021). A novel 3d laser scanning defects detection and measurement approach for automated fibre placement (afp) in-process inspection. *Measurement Science and Technology*, 32.

Zambal, S., Heindl, C., Eitzinger, C., & Scharinger, J. (2019). End-to-end defect detection in automated fiber placement based on artificially generated data. In Cudel, C., Bazeille, S., & Verrier, N. (Eds.), *Fourteenth International Conference on Quality Control by Artificial Vision*, volume 11172, (pp. 111721G). International Society for Optics and Photonics, SPIE.

Zemzemoglu, M. & Unel, M. (2022). Design and implementation of a vision based in-situ defect detection system of automated fiber placement process. In *2022 IEEE 20th International Conference on Industrial Informatics (INDIN)*, (pp. 393–398). IEEE.

Zemzemoglu, M. & Unel, M. (2023). A hierarchical learning-based approach for the automatic defect detection and classification of afp process using thermography. In *IECON 2023- 49th Annual Conference of the IEEE Industrial Electronics Society*, (pp. 1–6).

Zemzemoglu, M. & Unel, M. (2024). Thermal inspection and quality assessment for afp processes via automatic defect detection and segmentation. In *IECON 2024 - 50th Annual Conference of the IEEE Industrial Electronics Society*, (pp. 1–6).

Zemzemoglu, M., Unel, M., & Tunc, L. T. (2024). Enhancing automated fiber placement process monitoring and quality inspection: A hybrid thermal vision based framework. *Composites Part B: Engineering*, 285, 111753.

Zhang, Y., Wang, W., Liu, Q., Guo, Z., & Ji, Y. (2022). Research on defect detec-

tion in automated fiber placement processes based on a multi-scale detector. *Electronics*, 11(22), 3757.

Zhang, Z. (2000). A flexible new technique for camera calibration. *IEEE Transactions on pattern analysis and machine intelligence*, 22(11), 1330–1334.