Structural Representation Learning for Thermal Turbulence Detection in Infrared Imagery using YOLO

Akash Deep[1], Subhamoy Sen[1], and Arvind Keprate[2]

[1] *School of Civil and Environmental Engineering, Indian Institute of Technology Mandi, India*
d24221@students.iitmandi.ac.in
subhomay@iitmandi.ac.in

[2] *Green Energy Lab, Department of Mechanical, Electrical and Chemical Engineering, Oslo Metropolitan University, Norway*
arvind.keprate@oslomet.no

ABSTRACT

Thermal turbulence degrades imaging performance in long-range infrared systems by introducing spatially varying distortions that appear as irregular intensity fluctuations and curvilinear patterns. Detecting these regions is challenging due to the absence of well-defined boundaries and their diffuse nature. This work investigates how structural characteristics of thermal turbulence influence automated detection using deep learning–based object detectors. A systematic study is conducted to evaluate different structural representations derived from thermal imagery, including rolling guidance filter (RGF), variance-based fluctuation maps, curvature-based features from the Hessian matrix, and multi-scale vesselness responses using the Frangi filter. These descriptors are incorporated as multi-channel inputs within a YOLO-based detection framework and evaluated on annotated infrared turbulence data. Results show that while deep detectors can capture turbulence cues from raw thermal images, structural representations improve the visibility of distortions and enhance detection robustness. In addition, intensity-based enhancement strategies are analysed to examine whether simple contrast amplification alone can improve turbulence detection performance. A structural fusion of thermal intensity and complementary feature representations achieves the best overall performance, improving localisation accuracy and recall. The findings highlight the importance of representation design in detecting diffuse thermal patterns and provide a more reliable framework for turbulence-aware detection in infrared imagery.

1. INTRODUCTION

Infrared thermography has become an important non-contact sensing technique for monitoring the thermal behaviour of engineering systems and structures (Usamentiaga et al.,). Thermal imaging enables the observation of spatial temperature distributions that reflect underlying heat transfer processes (Fokaides & Kalogirou,). These temperature fields can provide valuable diagnostic information, as variations in thermal behaviour may indicate material defects, abnormal energy dissipation, or changes in structural performance (Sirca Jr & Adeli,). Because thermographic measurements can capture temperature variations across large surfaces without physical contact, infrared imaging is widely used in inspection and monitoring applications involving infrastructure, energy systems, and aerodynamic structures. In aerodynamic systems, surface temperature patterns are closely related to the interaction between airflow and structural surfaces. The characteristics of the boundary layer and the associated convective heat transfer processes strongly influence the thermal response of aerodynamic components. Changes in aerodynamic conditions such as boundary layer transition, flow separation, or surface degradation can alter local heat transfer behaviour and produce observable variations in surface temperature (Davis & Atkins,). Consequently, infrared thermography has been widely applied to study flow behaviour and aerodynamic performance, particularly in applications involving aircraft surfaces and wind turbine blades.

Recent investigations have shown that turbulent flow structures can generate distinctive thermal signatures in infrared imagery. These signatures appear as thermal turbulence patterns that arise from variations in convective heat transfer caused by turbulent airflow interacting with the surface (Carlomagno & Cardone,). Thermographic observations of wind turbine rotor blades, for example, have revealed that laminar–turbulent transition and flow separation can produce

visible thermal patterns along the blade surface. Such patterns reflect underlying aerodynamic conditions and provide useful information about airflow behaviour over the structure (Oehme et al.,). From a monitoring perspective, these thermal turbulence patterns are particularly important because aerodynamic degradation often begins with subtle changes in flow behaviour before significant structural damage occurs. Changes in surface roughness, contamination, erosion, or geometric deformation can modify boundary layer behaviour and influence convective heat transfer along aerodynamic surfaces. These changes may manifest in infrared imagery as altered turbulence patterns, which can therefore serve as indirect indicators of aerodynamic behaviour changes. Detecting and analysing these patterns may provide valuable insights for monitoring aerodynamic systems and identifying performance changes at an early stage (Feldmann et al.,).

Despite their diagnostic potential, identifying thermal turbulence patterns in infrared imagery remains challenging. Unlike conventional structural defects, turbulence patterns do not appear as discrete objects with well-defined boundaries. Instead, they typically manifest as irregular and spatially diffuse structures embedded within the background thermal field. Their appearance can vary depending on flow conditions, environmental factors, and surface properties, which complicates their reliable identification using traditional image analysis techniques (Chaudhuri et al.,). Recent developments in machine learning and computer vision have enabled automated analysis of complex patterns in imaging data. Deep learning based detection models have demonstrated strong performance in identifying spatial patterns across a variety of engineering and inspection tasks. These models offer a promising approach for automatically identifying turbulence patterns in thermographic imagery. However, most detection frameworks are designed to identify discrete objects, whereas turbulence patterns represent spatial variations in heat transfer rather than distinct physical entities. This highlights a fundamental challenge: conventional detection models rely on object-centric representations, whereas thermal turbulence is inherently a spatial and structural phenomenon, making its representation within image data a critical factor for reliable detection.

For this reason, the representation of turbulence patterns within image data becomes a critical factor for effective detection. Structural representations that emphasise turbulence-related features such as local thermal fluctuations, residual distortions, or curvilinear flow structures may help highlight characteristics associated with turbulent airflow, while also reflecting the inherent limitations of individual representations. Incorporating such representations into detection frameworks may therefore improve the ability of automated methods to identify turbulence patterns within infrared imagery (Banks,). As illustrated in the Figure 1, thermal turbulence patterns are inherently subtle and exhibit low contrast

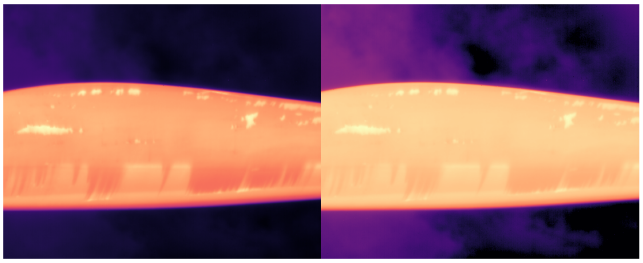

Figure 1. Original thermographic image (right) and a visually enhanced representation (left), illustrating the subtle and low-contrast nature of thermal turbulence patterns, which are not easily distinguishable in the raw thermal field.

against the background thermal field. In the original image, these structures are difficult to distinguish due to weak intensity variations. Contrast enhancement improves their visibility, highlighting the underlying distortions that are otherwise not easily detectable. While structural representations aim to explicitly encode turbulence-related spatial characteristics, it is also important to examine whether simpler intensity-based transformations can improve detection performance. In thermographic imagery, turbulence patterns are often subtle and embedded within low-contrast regions, suggesting that contrast enhancement techniques may improve their visibility. However, it remains unclear whether such intensity-based approaches are sufficient for reliable detection or whether explicit structural encoding is required.

Motivated by this, the present study systematically investigates how different structural representations of thermal turbulence influence detection performance. Several descriptors are explored and integrated into a deep learning–based detection framework to evaluate their effectiveness in identifying turbulence regions. The objective is to better understand the role of representation in turbulence detection and to provide a more reliable framework for analysing diffuse thermal patterns in infrared imagery. The remainder of the paper is as follows: section 2 contains background work, section 3 contains methodology opted in this work, section 4 contains results and discussions and section 5 contains conclusion and future scope.

2. RELATED WORK

Infrared thermography has been widely used for monitoring thermal behaviour in engineering systems, particularly in applications where heat transfer processes provide insight into system performance (Bagavathiappan, Lahiri, Saravanan, Philip, & Jayakumar,). In structural health monitoring and infrastructure inspection, thermographic methods have been employed to detect anomalies in materials and structures by observing spatial temperature variations on exposed surfaces. Because surface temperature distributions are influenced by internal heat transfer mechanisms, thermographic

imaging has been successfully applied to identify subsurface defects, insulation failures, and abnormal energy dissipation in structural components. In aerodynamic systems, thermal imaging has also been used to study flow behaviour, as surface temperature patterns are strongly affected by the interaction between airflow and structural surfaces (Sirca Jr & Adeli,).

The relationship between convective heat transfer and aerodynamic flow behaviour has been explored in several experimental studies. Variations in boundary layer characteristics influence the heat transfer between airflow and surfaces, producing measurable temperature patterns that can be observed using infrared imaging (Carlomagno & Cardone,). In particular, laminar–turbulent transition and flow separation phenomena are known to alter local heat transfer conditions, resulting in distinctive thermal signatures along aerodynamic surfaces. Such thermal features have been used to visualise boundary layer transitions and to assess aerodynamic performance in wind tunnel experiments and field observations (Davis & Atkins,).

Recent developments in thermographic imaging have enabled detailed observation of turbulence-related thermal patterns on operational aerodynamic structures. For example, Chaudhuri et al. introduced the KI-VISIR dataset, which contains thermographic imagery of turbulence patterns observed on wind turbine rotor blades during operation (Chaudhuri et al.,). In this dataset, turbulent flow regions appear as spatial thermal patterns associated with variations in convective heat transfer along the blade surface. These patterns reflect underlying aerodynamic flow structures and provide valuable information about flow transitions and aerodynamic conditions. The availability of annotated thermographic datasets has created new opportunities for analysing turbulence patterns using data-driven methods.

While thermographic imaging can reveal turbulence structures, atmospheric and optical effects may also influence the appearance of thermal patterns in captured imagery (Uzun & Akagündüz,). Atmospheric turbulence has long been recognised as a source of degradation in long-range imaging systems, as fluctuations in the refractive index of air can distort optical wave propagation and introduce geometric distortions and intensity fluctuations in images. Early investigations demonstrated that infrared imaging can capture observable turbulence structures through temporal and spatial variations in thermal imagery (Yasarla & Patel,). For instance, Watt and McHugh examined the use of infrared imaging for analysing atmospheric turbulence and estimating wind velocity through correlation analysis of sequential thermal images. These studies demonstrated that turbulence phenomena can be observed through imaging techniques and may exhibit structured spatial behaviour in thermal data (Watt & McHugh,).

A significant portion of the literature has focused on mitigating the impact of turbulence distortions in imaging systems. Classical image restoration approaches employ multi-frame processing techniques such as image registration, frame selection, and deconvolution to reconstruct undistorted images from turbulence-affected sequences(Lau, Lai, & Lui,). More recently, deep learning methods have been proposed to restore images degraded by atmospheric turbulence. For example, Yasarla and Patel introduced convolutional neural network architectures capable of learning turbulence-induced distortions and reconstructing cleaner images (Yasarla & Patel,). Although these approaches can improve image quality, they typically treat turbulence as a degradation artifact that should be removed rather than as a phenomenon that may contain useful structural information.

More recent studies have explored the application of machine learning models for detecting turbulence patterns in thermographic imagery. Object detection frameworks based on deep neural networks have been applied to datasets such as KI-VISIR to identify turbulence regions automatically. For instance, Ekici proposed an enhanced YOLO-based detection framework incorporating attention mechanisms and feature fusion strategies for detecting turbulence patterns in thermographic images of wind turbine blades(Ekici, Uyar, & Karadeniz,). These studies demonstrate the feasibility of automated turbulence detection using modern detection architectures. However, most existing approaches focus primarily on improving detection models rather than examining how turbulence patterns themselves should be represented in thermographic imagery. From an aerodynamic monitoring perspective, turbulence patterns observed in infrared images are not discrete objects but spatial thermal structures resulting from variations in convective heat transfer along aerodynamic surfaces. Their irregular and diffuse nature makes them difficult to represent using conventional object detection frameworks (Chaudhuri et al.,).

Beyond infrared thermography and aerodynamic flow analysis, data-driven monitoring frameworks have been increasingly adopted in the broader PHM community for real-time assessment of complex engineering systems. In offshore renewable energy systems, deep learning–based surrogate and indirect sensing approaches have been employed for fatigue assessment, damage detection, and condition monitoring of floating offshore wind turbine (FOWT) mooring systems using platform response information rather than direct instrumentation (Kumar, Sen, & Keprate, ,). These studies demonstrate the growing role of representation learning, temporal feature extraction, and data-driven inference for extracting physically meaningful monitoring information from complex operational signals(Arya, Kumar, Keprate, & Sen,)(Kumar, Thakur, Sen, & Keprate,). Although the present work focuses on thermal turbulence detection in infrared imagery, it shares a similar motivation of identifying informative repre-

sentations that enable reliable monitoring and early detection of performance-related phenomena.

Consequently, there remains a limited understanding of how different structural representations of turbulence patterns influence the ability of automated methods to identify them. In particular, descriptors that emphasise turbulence-related features such as thermal fluctuations, residual distortions, or curvilinear flow structures have not been systematically investigated in the context of thermographic turbulence detection. To address this limitation, the present study investigates multiple structural representations of thermal turbulence patterns and evaluates their effectiveness within a deep learning–based detection framework. By analysing how different representations influence turbulence detection performance, the study aims to improve the characterisation of thermal turbulence patterns and support their use as early indicators of aerodynamic degradation in infrared-based monitoring systems.

3. METHODOLOGY

3.1. Thermal Turbulence Pattern Formation

Thermal turbulence patterns (TTPs) observed in infrared imagery arise from variations in convective heat transfer between the aerodynamic surface and the surrounding airflow. In aerodynamic systems such as wind turbine blades, the surface temperature distribution is governed by the balance between internal heat conduction and external convective heat transfer.

The convective heat flux between the surface and the free-stream airflow can be expressed as

$$q = h(T_s - T_\infty) \tag{1}$$

where q denotes the convective heat flux, h represents the convective heat transfer coefficient, T_s is the surface temperature, and T_∞ denotes the free-stream air temperature.(Lienhard,)

The convective heat transfer coefficient h depends strongly on the flow regime. Under laminar flow conditions, the boundary layer remains stable and heat transfer occurs in a relatively uniform manner. When the flow transitions to turbulence, enhanced mixing within the boundary layer increases heat transfer rates, producing spatial variations in the surface temperature distribution. These variations manifest as irregular thermal patterns in infrared imagery (Incropera,).

Changes in aerodynamic performance, such as surface contamination, erosion, or geometric deformation, can modify boundary layer behaviour and induce local turbulence. Such changes alter the local heat transfer coefficient and consequently modify the thermal signature of the surface. Thermal turbulence patterns therefore provide indirect information about aerodynamic flow behaviour.

3.2. Infrared Thermographic Dataset

The analysis is conducted using thermographic imagery obtained from the KI-VISIR dataset. This dataset contains infrared images of operational wind turbine rotor blades captured under real atmospheric conditions. The images document thermal turbulence patterns generated by airflow interactions along the blade surface.

Each image captures spatial temperature variations across the blade surface, allowing visualisation of turbulence structures associated with laminar–turbulent transition and flow separation. The dataset also includes manually annotated regions corresponding to turbulence patterns, enabling supervised learning for automated detection.

These annotated regions are used as reference labels for evaluating the effectiveness of different structural representations in identifying turbulence patterns.

3.3. Thermal Image Preprocessing

Infrared images obtained from thermographic cameras often contain auxiliary information such as colour scales and measurement legends that are unrelated to the aerodynamic surface. These regions are removed during preprocessing to ensure that the analysis focuses only on the thermal field corresponding to the blade surface.

Following cropping, the thermal intensity values are normalised in order to reduce variations in temperature scale between different images. The normalised thermal intensity $I_n(x, y)$ is defined as

$$I_n(x, y) = \frac{I(x, y) - I_{min}}{I_{max} - I_{min}} \tag{2}$$

where $I(x, y)$ denotes the original thermal intensity at pixel location (x, y) and I_{min} and I_{max} represent the minimum and maximum intensity values within the image.(Raschka,)

Normalisation ensures consistent representation of thermal patterns across the dataset and improves the stability of subsequent feature extraction processes (Schwarzhans et al.,).

3.4. Statistical Analysis of Thermal Fluctuations

To further support the physical interpretation of thermal turbulence patterns, a statistical analysis of spatial temperature fluctuations was conducted using the KI-VISIR thermographic dataset. Only images containing valid turbulence annotations were considered, resulting in a curated subset of 689 samples.

For each thermographic frame, the spatial temperature field

was analysed by computing the local variance of the normalised thermal intensity using a sliding neighbourhood window. The local variance at each pixel location (x, y) is defined as

$$\sigma^2(x,y) = \frac{1}{|\Omega|} \sum_{(i,j)\in\Omega} \left(I_n(i,j) - \mu(x,y)\right)^2 \tag{3}$$

where Ω denotes the neighbourhood window and $\mu(x, y)$ represents the local mean intensity within that neighbourhood.

Using the polygon annotations provided in the dataset, variance samples were extracted separately from annotated turbulence regions and from surrounding blade areas representing the background thermal field. For each image, balanced random sampling was applied to avoid bias due to differing region sizes.

Figure 2 illustrates the statistical distribution of local temperature variance for turbulence and background regions across the analysed dataset. The results indicate that turbulence regions exhibit consistently higher spatial variance compared to the surrounding blade surface.

This observation is consistent with the physical behaviour of turbulent flow. When the boundary layer transitions to turbulence, enhanced mixing increases fluctuations in convective heat transfer, producing stronger spatial variations in the surface temperature field. Consequently, turbulence structures appear in thermographic imagery as regions of elevated temperature variance.

A statistical comparison using Welch's t-test confirms that the difference between the two distributions is highly significant ($p < 0.001$) with a large effect size (Cohen's $d > 0.8$). These findings provide quantitative evidence that turbulence regions are characterised by stronger thermal fluctuations, supporting the use of variance-based descriptors for highlighting turbulence structures in infrared imagery.

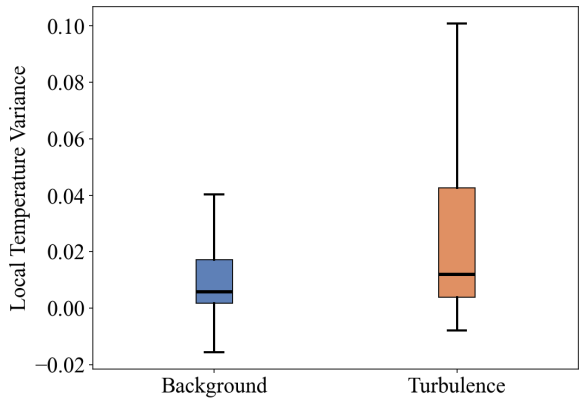

Figure 2. Distribution of local temperature variance in turbulence and background regions.

3.5. Gradient-Based Analysis of Thermal Transitions

To further investigate the spatial characteristics of thermal turbulence patterns, the gradient magnitude of the temperature field was analysed. The temperature gradient provides a measure of how rapidly temperature changes across the blade surface and therefore highlights regions with strong thermal transitions.

The gradient magnitude was computed using Sobel operators to estimate the first-order spatial derivatives of the thermal intensity field

$$|\nabla T(x,y)| = \sqrt{\left(\frac{\partial T}{\partial x}\right)^2 + \left(\frac{\partial T}{\partial y}\right)^2} \tag{4}$$

where $\frac{\partial T}{\partial x}$ and $\frac{\partial T}{\partial y}$ denote the spatial derivatives of the temperature field along the horizontal and vertical directions.

Using the annotated turbulence regions provided in the KI-VISIR dataset, gradient magnitude values were extracted from turbulence regions and from surrounding blade background. Figure 3 illustrates the distribution of gradient magnitudes for the two regions.

The results indicate that turbulence regions exhibit larger temperature gradients and greater variability compared to the background thermal field. This behaviour is consistent with the physics of turbulent boundary layers, where enhanced mixing produces sharper spatial variations in convective heat transfer. A statistical comparison using Welch's t-test confirms that the difference between the two distributions is significant ($p < 0.01$).

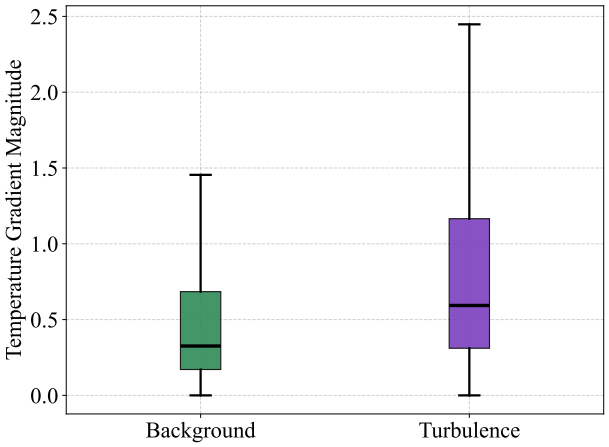

Figure 3. Distribution of temperature gradient magnitude in turbulence and background regions.

3.6. Structural Representation of Thermal Turbulence Patterns

Thermal turbulence patterns appear as irregular spatial structures embedded within the background temperature field. Unlike conventional defects, these patterns do not correspond to discrete objects with well-defined boundaries. Instead, they manifest as local distortions, fluctuations, and curvilinear structures associated with variations in convective heat transfer.

To capture these characteristics, several structural representations are explored to emphasise turbulence-related features within the thermal imagery.

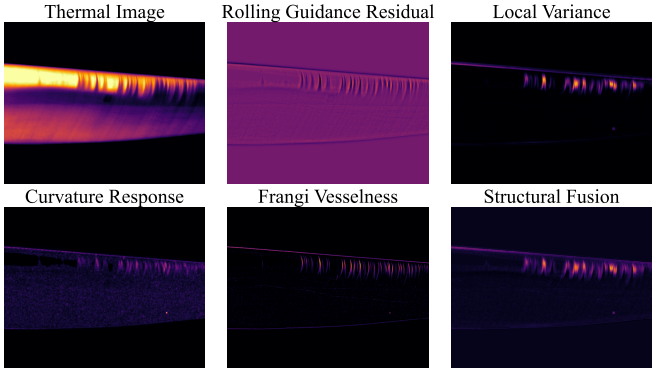

Figure 4. Structural representations of thermal turbulence patterns in infrared imagery of a wind turbine blade, showing residual, variance, curvature, and vesselness responses used to highlight turbulence structures for detection.

3.6.1. Rolling Guidance filter Representation

The rolling guidance filter (Zhang, Shen, Xu, & Jia,) is used to separate large-scale thermal variations from small-scale distortions associated with turbulence. The filter performs iterative edge-preserving smoothing using a bilateral filtering operation.

The filtered image $S(x, y)$ at iteration k is defined as

$$S^{k+1}(x,y) = \frac{1}{W(x,y)} \sum_{(i,j)\in\Omega} I(i,j) \exp\left(-\frac{\|(x,y)-(i,j)\|^2}{2\sigma_s^2}\right) \exp\left(-\frac{|S^k(x,y)-S^k(i,j)|^2}{2\sigma_r^2}\right) \quad (5)$$

where Ω represents the neighbourhood window, σ_s controls spatial smoothing, σ_r controls range similarity, and $W(x, y)$ is the normalisation factor.

The residual representation highlighting small-scale thermal distortions is then obtained as

$$R(x,y) = I_n(x,y) - S(x,y) \quad (6)$$

This residual component emphasises fine-scale thermal variations associated with turbulence structures.

3.6.2. Local Variance Representation

Turbulent regions tend to exhibit stronger spatial fluctuations in temperature due to variations in convective heat transfer. These fluctuations can be captured using a local variance representation (Haralick, Shanmugam, & Dinstein,).

For a neighbourhood window Ω, the local variance at pixel location (x, y) is computed as

$$\sigma^2(x,y) = \frac{1}{|\Omega|} \sum_{(i,j)\in\Omega} (I_n(i,j) - \mu(x,y))^2 \quad (7)$$

where $\mu(x, y)$ represents the local mean intensity defined as

$$\mu(x,y) = \frac{1}{|\Omega|} \sum_{(i,j)\in\Omega} I_n(i,j) \quad (8)$$

Regions with elevated variance correspond to areas of strong thermal fluctuation and are therefore indicative of turbulence activity.

3.6.3. Curvature-Based Representation

Turbulence structures frequently exhibit curvilinear spatial patterns associated with flow instabilities. These structures can be emphasised by analysing second-order derivatives of the thermal intensity field.

The local curvature properties of the image are described using the Hessian matrix

$$H(x,y) = \begin{bmatrix} I_{xx}(x,y) & I_{xy}(x,y) \\ I_{xy}(x,y) & I_{yy}(x,y) \end{bmatrix} \quad (9)$$

where I_{xx}, I_{yy}, and I_{xy} denote second-order partial derivatives of the image intensity.

The eigenvalues of the Hessian matrix provide information about local curvature characteristics and highlight filament-like structures present in the thermal image (Yang et al.,).

3.6.4. Frangi Vesselness Representation

To further enhance curvilinear turbulence structures, a multi-scale vesselness filter based on the Hessian matrix is applied. The vesselness response is computed as

$$V(x,y) = \exp\left(-\frac{R_B^2}{2\beta^2}\right)\left(1-\exp\left(-\frac{S^2}{2c^2}\right)\right) \tag{10}$$

where

$$R_B = \frac{|\lambda_1|}{|\lambda_2|} \tag{11}$$

and

$$S = \sqrt{\lambda_1^2 + \lambda_2^2} \tag{12}$$

with λ_1 and λ_2 representing the eigenvalues of the Hessian matrix. Parameters β and c control the sensitivity of the filter to elongated structures (Frangi, Niessen, Vincken, & Viergever,).

This representation highlights filament-like thermal features corresponding to turbulence structures.

3.7. Structural Fusion Representation

The structural fusion representation integrates complementary physical characteristics of turbulence by combining thermal intensity, residual distortions, and structural flow features.

$$X = [I_n, R, C] \tag{13}$$

where I_n represents the normalised thermal intensity, R captures high-frequency turbulence-induced distortions, and C encodes curvature-based structural flow patterns.

Although variance-based descriptors effectively highlight local thermal fluctuations, they were not incorporated into the final structural fusion representation due to their partial redundancy with residual-based features. The rolling guidance filter representation already captures fine-scale intensity variations and turbulence-induced local distortions, which overlap substantially with fluctuation information encoded by variance maps. Preliminary comparative analysis indicated that including variance descriptors within the fusion framework did not produce a consistent improvement in localisation performance while increasing feature redundancy and input complexity. Consequently, the final fusion formulation prioritises complementary representations that encode thermal intensity, high-frequency residual behaviour, and structural flow characteristics, enabling a more compact and physically interpretable representation of turbulence patterns.

This fusion enables the model to jointly capture thermal intensity variations, local turbulence-induced distortions, and structural flow continuity. By combining complementary representations, the structural fusion approach provides a more informative input space for learning turbulence characteristics, leading to improved localisation performance.

3.8. Intensity-Based Representation Analysis

In addition to structural representations, intensity-based transformations were evaluated to investigate whether enhancing thermal contrast alone can improve turbulence detection. The normalised thermal intensity field $I_n(x,y)$ is first obtained as defined in Eq. (2).

Contrast enhancement is applied using adaptive histogram equalisation, which locally redistributes intensity values to improve visibility of subtle patterns. The enhanced intensity field $I_c(x,y)$ can be expressed as

$$I_c(x,y) = \mathcal{H}\left(I_n(x,y)\right) \tag{14}$$

where $\mathcal{H}(\cdot)$ denotes the adaptive histogram equalisation operator.

Unlike structural representations, this transformation does not explicitly encode spatial derivatives or geometric characteristics of turbulence patterns, but instead modifies the local intensity distribution.

To analyse the interaction between intensity and structural information, hybrid representations were constructed by combining contrast-enhanced intensity with structural descriptors such as residual and curvature features. For example, a contrast–residual hybrid representation can be expressed as

$$X_{CR}(x,y) = [I_c(x,y), R(x,y)] \tag{15}$$

and a contrast–residual and curvature representation is defined as

$$X_{CRC}(x,y) = [I_c(x,y), R(x,y), C(x,y)] \tag{16}$$

where $R(x,y)$ represents the residual component defined in Eq. (6), and $C(x,y)$ denotes the curvature-based representation derived from second-order derivatives as defined in Eq. (9).

These formulations allow the model to jointly learn from intensity-enhanced features and structural characteristics, enabling a systematic comparison between intensity-based and structure-based representations.

4. RESULTS AND DISCUSSION

The proposed framework was evaluated using thermographic images from the KI-VISIR dataset in order to analyse how different structural representations of thermal turbulence patterns influence detection performance. The experiments were

designed in two stages. First, multiple object detection architectures were evaluated on the raw thermographic dataset to identify the most suitable detector for turbulence localisation. Subsequently, the selected detector was used to investigate the effect of different structural representations derived from the thermal field. The dataset was divided into training, validation, and testing subsets using a consistent split across all experiments to ensure fair comparison. The experiments were conducted on a curated subset of the KI-VISIR dataset consisting of 689 thermographic images with annotated turbulence regions. The dataset was divided into training(80%), validation(10%), and testing(10%) sets using a consistent split across all experiments. Annotations are provided in the form of bounding boxes corresponding to turbulence regions, enabling supervised detection training and evaluation. Model performance was evaluated using standard object detection metrics, including precision (P), recall (R), mean Average Precision at an IoU threshold of 0.5 (mAP_{50}), and mean Average Precision across IoU thresholds from 0.5 to 0.95 (mAP_{50-95}). These metrics provide complementary insight into detection accuracy as well as localisation quality.

4.1. Experimental Setup

All detection experiments were implemented using the Ultralytics YOLO framework to maintain a consistent training and evaluation pipeline across model architectures and representation types. The analysed KI-VISIR subset consisting of 689 annotated thermographic images was divided into training (80%), validation (10%), and testing (10%) sets using an identical split across all experiments to ensure fair comparison. All models were trained for 100 epochs using an input resolution of 640×640 pixels. Optimisation was performed using the AdamW optimiser, while consistent training settings, augmentation procedures, and inference criteria were maintained throughout all experiments. Confidence thresholding and Intersection-over-Union (IoU) filtering were applied uniformly during evaluation to ensure comparable localisation assessment. To minimise variability arising from stochastic training behaviour, identical framework settings and data partitioning strategies were preserved across experiments. This controlled experimental setup ensures that observed performance differences primarily reflect the influence of the investigated structural representations rather than variations in optimisation, implementation, or dataset allocation.

4.2. Detector Performance on the Raw Dataset

The first experiment evaluates the ability of different YOLO detection architectures to identify thermal turbulence patterns directly from the raw thermographic images. Table 1 summarises the detection performance obtained for YOLOv8-s, YOLOv9-s, YOLOv10-s, and YOLOv11-s. Among the evaluated models, YOLOv11-s achieves the highest localisation performance with an mAP_{50} of 0.872 and an mAP_{50-95} of 0.591. This indicates superior bounding box alignment and localisation quality for turbulence structures. Although YOLOv8-s and YOLOv9-s demonstrate competitive performance, their localisation accuracy is slightly lower. YOLOv10-s shows comparatively reduced performance, likely due to differences in feature aggregation strategies.Based on these results, YOLOv11-s is selected as the primary detector due to its superior localisation accuracy and balanced precision–recall performance.

Table 1. Detection performance of YOLO models on raw thermographic dataset.

Model	P	R	mAP_{50}	mAP_{50-95}
YOLOv8-s	0.803	0.774	0.862	0.584
YOLOv9-s	0.802	0.775	0.861	0.579
YOLOv10-s	0.771	0.786	0.851	0.573
YOLOv11-s	0.819	0.780	0.874	0.589

4.3. Detection Performance Across Structural Representations

After selecting YOLOv11-s as the baseline detector, the next set of experiments investigates how different structural representations of the thermal field influence turbulence detection performance. The evaluated representations include rolling guidance filter maps, local variance maps, curvature responses derived from second-order derivatives, and Frangi vesselness filters that highlight curvilinear turbulence structures.Table 2 summarises the detection performance obtained using these different structural representations.

Table 2. Performance of structural representations using YOLOv11.

Representation	P	R	mAP_{50}	mAP_{50-95}
Raw	0.819	0.780	0.874	0.589
RGF	0.792	0.802	0.869	0.575
Frangi	0.823	0.759	0.865	0.572
Curvature	0.797	0.765	0.854	0.550
Structural fusion	0.789	0.813	0.877	0.602

The results indicate that the raw thermographic representation already provides strong detection performance, suggesting that turbulence patterns are inherently encoded in the thermal field. Individual structural representations improve interpretability by emphasising specific turbulence characteristics such as residual distortions and curvilinear structures. However, these representations do not consistently improve detection performance.

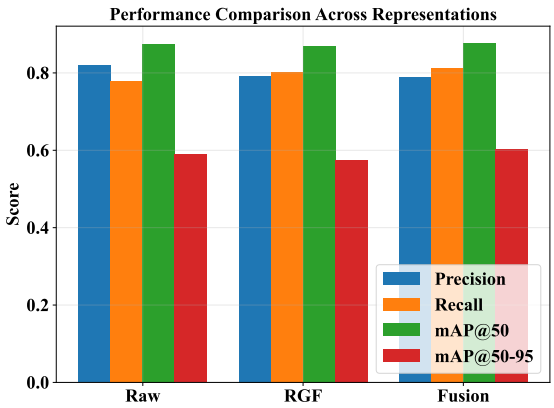

Figure 5. Comparison of detection performance across structural representations (top 3). Structural fusion achieves the best overall balance, improving recall and localisation accuracy (mAP_{50-95}) while maintaining competitive precision.

The structural fusion representation achieves the best overall performance, improving mAP_{50} to 0.877 and mAP_{50-95} to 0.602 while also increasing recall to 0.813. This demonstrates that combining complementary turbulence descriptors enhances localisation accuracy. This also suggests that no single representation fully captures the multi-scale and structural complexity of turbulence patterns, whereas their combination provides a more complete description of the underlying thermal field.

4.4. Structural Fusion Representation

To further exploit complementary information contained in different structural descriptors, a structural fusion representation was constructed by combining the thermal intensity, residual, and curvature-based representations. This representation integrates information related to thermal distortions, statistical fluctuations, and curvilinear turbulence structures within a unified feature space.

The structural fusion representation achieves the best overall balance between recall and localisation performance. Although the overall mAP values remain comparable to the baseline, the fusion representation provides a more interpretable description of turbulence structures within the thermographic imagery. By integrating complementary descriptors, the fused representation highlights regions where turbulence-induced thermal distortions are most pronounced, improving the reliability of turbulence localisation. Figure 5 provides a comparative view of detection performance across representations. While the raw thermal input already achieves strong results, individual structural descriptors introduce trade-offs between precision and recall. The structural fusion representation consistently achieves the best balance, indicating that combining complementary features improves both detection reliability and localisation quality. Overall, these results demonstrate that structural representations enhance detection performance by improving feature interpretability and localisation quality, with structural fusion providing the most balanced improvement across evaluation metrics.

Table 3. Performance comparison of intensity-enhanced and structural representations using YOLOv11. Results are reported using consistent three-decimal precision for fair comparison across representations.

Representation	P	R	mAP_{50}	mAP_{50-95}
Raw	0.819	0.780	0.874	0.589
CL + Raw	0.790	0.780	0.850	0.550
CL + RGF	**0.820**	0.770	0.870	0.570
CL + Frangi	0.810	0.720	0.840	0.560
CL + Curvature	0.750	**0.820**	0.860	0.550
Structural Fusion	0.789	0.813	**0.877**	**0.602**

Although the structural fusion representation achieves the strongest overall balance across evaluation metrics, the observed performance improvements remain relatively moderate. However, even incremental gains are meaningful in the context of thermal turbulence detection, as improved recall and localisation accuracy directly increase the number of cases in which turbulence regions are correctly identified within complex thermographic fields. Since turbulence patterns are spatially diffuse, low-contrast, and often difficult to localise reliably, modest improvements can translate into a greater number of successfully detected turbulence instances, thereby enhancing monitoring sensitivity. Furthermore, the present analysis is conducted using a curated subset of the KI-VISIR dataset corresponding to a single application domain involving wind turbine blade thermography. Consequently, the reported gains should be interpreted within the scope of the analysed dataset and experimental configuration. While identical training settings and data partitioning were maintained across experiments to enable fair comparison, repeated split analysis, formal statistical significance testing, and validation across additional thermographic datasets were not considered within the present study. These directions represent important extensions for future work toward establishing broader generalisability and statistical robustness of representation-driven turbulence detection.

4.5. Effect of Intensity-Based Representations

To further investigate the role of intensity information, contrast-enhanced representations were evaluated using adaptive histogram equalisation and their combinations with structural descriptors. Table 3 summarises the results.

The results demonstrate that contrast enhancement alone (CLAHE (CL) + Raw) does not consistently improve detection performance over the baseline, indicating that turbu-

lence patterns are not solely limited by low contrast. When combined with structural representations, distinct behaviours emerge: CL + RGF improves precision, while CL + Curvature increases recall, highlighting a trade-off between localisation accuracy and detection sensitivity. The degradation observed with CL + Frangi suggests that turbulence patterns do not conform well to vessel-like assumptions. Overall, the structural fusion representation achieves the best balance across all metrics, confirming that effective turbulence detection requires complementary structural encoding rather than intensity enhancement alone.

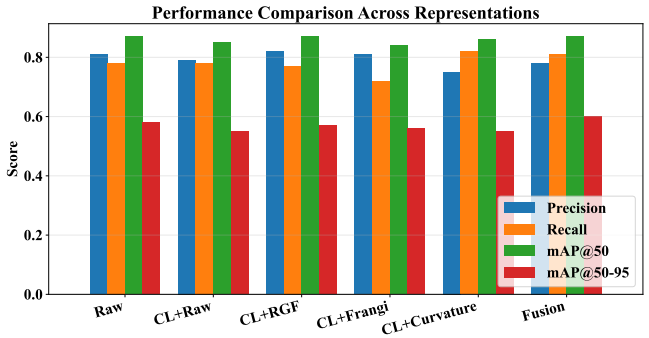

Figure 6. Performance comparison across intensity and structural representations. The CL (CLAHE)-based representations highlight the influence of contrast enhancement when combined with structural descriptors such as residual and curvature features.

Figure 6 illustrates the impact of intensity-based enhancement in combination with structural representations on detection performance. The CL + Raw representation shows that contrast enhancement alone does not provide consistent improvements over the baseline. When combined with structural descriptors, distinct behaviours emerge: CL + RGF improves precision by suppressing false detections, while CL + Curvature increases recall by capturing a broader range of turbulence patterns at the cost of reduced precision. The structural fusion representation achieves the best overall balance, confirming that effective turbulence detection depends on complementary structural encoding rather than intensity enhancement alone.

5. CONCLUSION AND FUTURE WORK

This study investigated the role of structural representation in detecting thermal turbulence patterns from infrared imagery. Unlike conventional object detection tasks, turbulence patterns represent spatially diffuse thermal structures without well-defined boundaries, making their representation a critical factor for reliable detection. A systematic evaluation of multiple structural representations, including rolling guidance filter residuals, curvature-based features, and vesselness responses, demonstrated that while individual descriptors enhance interpretability, they do not consistently improve detection performance. In contrast, the proposed structural fusion representation, which integrates thermal intensity, residual distortions, and structural flow characteristics, achieves the best overall performance, improving both localisation accuracy and recall.

In addition, intensity-based enhancement experiments using contrast-limited adaptive histogram equalisation (CLAHE) indicate that improving thermal contrast alone does not consistently enhance detection performance. While contrast-enhanced representations improve visual visibility, they fail to provide reliable gains without complementary structural encoding. The results demonstrate that structural fusion improves detection performance by combining complementary turbulence descriptors, leading to enhanced localisation accuracy and recall. These findings highlight the importance of representation design in detecting diffuse thermal patterns and improving the reliability of turbulence-aware detection systems.Several directions can be explored to further extend this work. First, Grad-CAM–based interpretability analysis can be incorporated to verify that model attention aligns with physically meaningful turbulence regions. Second, temporal consistency analysis using sequential thermographic data can provide deeper insight into the evolution of turbulence structures over time. Third, the integration of structural constraints or flow-based priors into deep learning models may further improve robustness and generalisation. Finally, extending the framework to real-time monitoring systems for aerodynamic structures such as wind turbines can enable practical deployment for early detection of performance degradation.

From a practical deployment perspective, the proposed representation framework is relevant for infrared-based monitoring systems operating on aerodynamic structures such as wind turbine blades. Although the investigated structural representations introduce an additional preprocessing stage, the employed operations primarily involve image-based filtering and derivative computations that remain computationally lightweight compared to detector training and inference. Consequently, the framework has potential applicability in near real-time monitoring scenarios where improved turbulence localisation may support earlier identification of aerodynamic performance changes. In such settings, even moderate improvements in detection reliability may be operationally valuable, as they increase the number of correctly identified turbulence regions available for downstream condition assessment. A detailed computational analysis involving inference latency, preprocessing overhead, and hardware-dependent deployment performance represents an important direction for future work.

Overall, this study demonstrates that effective structural representation learning provides a more reliable and interpretable framework for turbulence-aware detection systems.

ACKNOWLEDGEMENTS

This work is funded by the Science and Engineering Research Board (SERB) scheme of the Department of Science and Technology (DST), Government of India, through grant file number SERB/CRG/2023/005267. The authors are sincerely grateful for the financial support and encouragement provided by the DST.

REFERENCES

Arya, A., Kumar, R., Keprate, A., Sen, S. (2025, Nov.). Real-time monitoring of clump weight integrity loss in floating wind turbines via deep learning. *Journal of Dynamics, Monitoring and Diagnostics*, *5*(1), 38–48. doi: 10.37965/jdmd.2025.792

Bagavathiappan, S., Lahiri, B. B., Saravanan, T., Philip, J., Jayakumar, T. (2013). Infrared thermography for condition monitoring–a review. *Infrared Physics & Technology*, *60*, 35–55.

Banks, D. (2000). *Visualization of in-flight flow phenomena using infrared thermography*. NASA Dryden Flight Research Center.

Carlomagno, G. M., Cardone, G. (2010). Infrared thermography for convective heat transfer measurements. *Experiments in fluids*, *49*(6), 1187–1218.

Chaudhuri, S., Stamm, M., Lapšanská, I., Lançon, T., Osterbrink, L., Driebe, T., ... Harendt, R. (2025). Infrared thermography of turbulence patterns of operational wind turbine rotor blades supported with high-resolution photography: Ki-visir dataset. *Wind Energy*, *28*(1), e2958.

Davis, W., Atkins, N. R. (2024). Infrared thermography techniques for boundary layer state visualisation. *Experiments in Fluids*, *65*(6), 91.

Ekici, S., Uyar, M., Karadeniz, T. N. (2025). A-biyolov9: An attention-guided yolov9 model for infrared-based wind turbine inspection. *Applied Sciences*, *15*(21), 11840.

Feldmann, D., Oehme, F., von Germersheim, L., López Parras, R., Fischer, A., Avila, M. (2022). Towards indirect assessment of surface anomalies on wind turbine rotor blades. In *Stab/dglr symposium* (pp. 529–538).

Fokaides, P. A., Kalogirou, S. A. (2011). Application of infrared thermography for the determination of the overall heat transfer coefficient (u-value) in building envelopes. *Applied energy*, *88*(12), 4358–4365.

Frangi, A. F., Niessen, W. J., Vincken, K. L., Viergever, M. A. (1998). Multiscale vessel enhancement filtering. In *International conference on medical image computing and computer-assisted intervention* (pp. 130–137).

Haralick, R. M., Shanmugam, K., Dinstein, I. H. (2007). Textural features for image classification. *IEEE Transactions on systems, man, and cybernetics*(6), 610–621.

Incropera, F. P. (2007). *Fundamentals of heat and mass transfer: Frank p. incropera...[et al.]*. John Wiley.

Kumar, R., Sen, S., Keprate, A. (2025). Real-time fatigue assessment of floating offshore wind turbine mooring employing sequence-to-sequence-based deep learning on indirect fatigue response. *Ocean Engineering*, *315*, 119741. doi: https://doi.org/10.1016/j.oceaneng.2024.119741

Kumar, R., Sen, S., Keprate, A. (2026). Fatigue damage assessment of fowt mooring lines using sequence-to-sequence based indirect sensing. In M. Singh et al. (Eds.), *Proceedings of the unified conference of damas, income viii and tepen conferences* (pp. 47–57). Cham: Springer Nature Switzerland.

Kumar, R., Thakur, A., Sen, S., Keprate, A. (2026). Temporal feature extraction based real time damage detection of floating offshore wind turbine mooring lines. In M. Singh et al. (Eds.), *Proceedings of the unified conference of damas, income viii and tepen conferences* (pp. 59–68). Cham: Springer Nature Switzerland.

Lau, C. P., Lai, Y. H., Lui, L. M. (2019). Restoration of atmospheric turbulence-distorted images via rpca and quasiconformal maps. *Inverse Problems*, *35*(7), 074002.

Lienhard, J. H. (2005). *A heat transfer textbook*. Phlogistron.

Oehme, F., Gleichauf, D., Suhr, J., Balaresque, N., Sorg, M., Fischer, A. (2022). Thermographic detection of turbulent flow separation on rotor blades of wind turbines in operation. *Journal of Wind Engineering and Industrial Aerodynamics*, *226*, 105025.

Raschka, S. (2014). About feature scaling and normalization and the effect of standardization for machine learning algorithms. *Political Leg Anthropology Rev*, *30*, 67–89.

Schwarzhans, F., George, G., Sanchez, L. E., Zaric, O., Abraham, J. E., Woitek, R., Hatamikia, S. (2025). Image normalization techniques and their effect on the robustness and predictive power of breast mri radiomics. *European Journal of Radiology*, *187*, 112086.

Sirca Jr, G. F., Adeli, H. (2018). Infrared thermography for detecting defects in concrete structures. *Journal of Civil Engineering and Management*, *24*(7), 508–515.

Usamentiaga, R., Venegas, P., Guerediaga, J., Vega, L., Molleda, J., Bulnes, F. G. (2014). Infrared thermography for temperature measurement and non-destructive testing. *Sensors*, *14*(7), 12305–12348.

Uzun, E., Akagündüz, E. (2025). Augmenting atmospheric turbulence effects on thermal-adapted deep object detection models. *Scientific Reports*, *15*(1), 9900.

Watt, D., McHugh, J. (1990). Infrared thermal imaging of atmospheric turbulence. In *Nasa, langley research center, airborne wind shear detection and warning systems. second combined manufacturers' and technologists' conference, part 1.*

Yang, J., Ma, S., Sun, Q., Tan, W., Xu, M., Chen, N., Zhao, D. (2014). Improved hessian multiscale enhancement

filter. *Bio-medical materials and engineering*, *24*(6), 3267–3275.

Yasarla, R., Patel, V. M. (2020). Learning to restore a single face image degraded by atmospheric turbulence using cnns. *arXiv preprint arXiv:2007.08404*.

Yasarla, R., Patel, V. M. (2021). Learning to restore images degraded by atmospheric turbulence using uncertainty. In *2021 ieee international conference on image processing (icip)* (p. 1694-1698). doi: 10.1109/ICIP42928.2021.9506614

Zhang, Q., Shen, X., Xu, L., Jia, J. (2014). Rolling guidance filter. In *European conference on computer vision* (pp. 815–830).