

Article

Efficient Reconstruction of High-Resolution Tidal Turbine Blade Deflection and Strain Maps Through Sensing Location Optimisation

Marek J. Munko ^{1,*} , Miguel A. Valdivia Camacho ¹ , Fergus Cuthill ¹ , Conchúr M. Ó Brádaigh ² 
and Sergio Lopez Dubon ^{1,*} 

¹ School of Engineering, University of Edinburgh, Mayfield Road, Edinburgh EH9 3JL, UK

² Faculty of Engineering, University of Sheffield, Sir Frederick Mappin Building, Mappin Street, Sheffield S1 3JD, UK

* Correspondence: m.munko@ed.ac.uk (M.J.M.); sergio.ldubon@ed.ac.uk (S.L.D.)

Abstract

During fatigue tests of tidal turbine blades, digital image correlation (DIC) is used to collect vital information about the specimen. DIC provides high-resolution displacement and strain maps of selected blade sections; however, continuous operation is hindered by the need to acquire, transfer, and process large volumes of high-resolution images, precluding real-time use during long tests. We address this problem by optimising sparse sensing locations on the blade surface so that full-field maps can be accurately reconstructed from a small subset of pixel measurements. In contrast to most DIC improvements found in the literature, which focus on accelerating the processing stage, this approach circumvents the need to collect high-resolution data. We evaluate this approach in a case study at FastBlade, a dedicated testing facility for tidal turbine blades. With less than 1% of the original pixels measured, the mean relative error evaluated on the dataset is 0.4% and 16% for displacement and strain maps, respectively, with the larger strain error reflecting the higher spatial complexity of strain fields. The optimised layouts outperform random and grid-like arrangements. The framework enables real-time monitoring and, subject to relevant validation, might be applied to reconstruct high-resolution strain maps directly from strain-gauge readings, potentially extending to in-ocean blade monitoring. Given the high accuracy of deflection reconstructions, using them to derive strain fields is suggested as a direction for further study.

Keywords: sensor placement optimisation; tidal turbine blade; fatigue testing; data-driven



Academic Editors: Joško Parunov and
Yordan Garbatov

Received: 12 January 2026

Revised: 11 February 2026

Accepted: 17 February 2026

Published: 24 February 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article
distributed under the terms and

conditions of the [Creative Commons
Attribution \(CC BY\) license](https://creativecommons.org/licenses/by/4.0/).

1. Introduction

Over the last few years, the energy market has seen a rapid increase in the net supply of electricity generated from renewable energy sources [1]. Although tidal energy is currently exploited to a significantly lesser extent than, for example, wind or solar energy, it is characterised by much better predictability [2]. Tidal turbine blades are among the key load-bearing components of tidal turbines, and their reliability significantly impacts the success of electricity generation. The testing procedure, which aims to de-risk the design and fabrication of the blade before its subsea deployment, relies extensively on a variety of sensor readings collected as the blade undergoes cyclic stressing. Some of the most valuable sources of information are the high-resolution displacement and strain maps collected with digital image correlation (DIC). However, the DIC process is computationally expensive,

prohibiting real-time operation over prolonged time intervals, and requires substantial data storage resources. Moreover, due to the nature of the setup, DIC data cannot be collected when the blade operates in the ocean. This work aims to investigate how optimising point sensing locations on a blade can help alleviate the computational burden, enabling real-time DIC operation while maintaining high-resolution information about the strain and displacement distribution.

1.1. Tidal Energy

The high predictability of tidal energy is the key advantage setting tidal energy apart from other renewables, such as wind or solar energy, which are subject to inherent variations [2]. Tidal energy has also proved to be abundant in many areas worldwide, such as the United Kingdom (UK), Canada, France, Norway, and Spain [3]. In the coming years, further development of tidal energy technology is expected to lower the levelised cost of energy (LCoE) and result in an increased contribution of this resource to the energy mix. Research shows that 30.2 MW of tidal energy capacity was already installed in Europe between 2010 and 2021 [4]. In the UK alone, electricity generation from tidal energy is projected to reach up to 11.5 GW [5], constituting 11% of the UK's energy demand [6].

1.2. Tidal Turbine Blade Testing

Given the harsh operating conditions of tidal turbines and their reduced accessibility during deployment, de-risking their design is particularly important to help the technology become more widespread. Blade tests involve subjecting it to cyclic loads, aiming to replicate the damage caused by a lifetime of ocean loading. Unlike wind turbine blades, which can be excited at their natural frequency to generate loading [7,8], composite tidal turbine blades have a higher natural frequency (around 18–20 Hz [9]). This high natural frequency does not allow for fatigue testing based on the natural frequency oscillation, requiring the mechanical introduction of loads. However, actuating a stiff blade results in higher costs due to the amount of energy needed to displace the specimen.

FastBlade

FastBlade, pictured in Figure 1a, is a research facility at the University of Edinburgh where slender structures of up to 20 m in length can be structurally tested. FastBlade is also the world's first dedicated facility for structural testing of tidal turbine blades, which is capable of recovering energy while operating at a frequency below the natural frequency of a specimen [10]. This makes the facility particularly relevant for testing stiff structures, thereby accelerating the development of tidal turbine blade technology. Among FastBlade's data collection systems, DIC provides non-intrusive measurements of the structure's surface deformation. Running DIC is computationally expensive, however, making it unfeasible to run in real time, and the vast amounts of data collected significantly impact the data storage capacity available at FastBlade.




Figure 1. (a) The test setup at FastBlade. Three hydraulic actuators are used to load the blade. During a test, they cyclically extend and contract, applying varying loads to the blade. (b) Left: The side view of the test setup at FastBlade. Right: The bottom view of the test setup. 1: The area of the blade examined, covered with the speckle pattern; 2: one of the saddles connecting the actuators to the blade; 3: one of the three actuators; 4: the strong wall on which the blade is mounted; 5: one of the lamps improving the consistency of the lighting conditions for DIC; 6: one of the four cameras used for DIC.

1.3. Paper Contributions

Specific contributions of this work include:

- Near-real-time, high-resolution reconstructions using data collected at FastBlade to show that measuring $<1\%$ of original pixels enables near-real-time DIC reconstruction, while preserving high spatial fidelity (yielding $\approx 0.4\%$ mean relative error for displacement and $\approx 16\%$ for strain).
- Benchmarking against baselines, demonstrating the superiority over random and grid-like sensor layouts using the FastBlade case study.
- Data volume reduction, eliminating the need to collect high-resolution raw images required by neural-network DIC-acceleration approaches [11,12], resulting in the reduction of data storage by over 99%.
- A step towards high-resolution in situ sensing, demonstrating a pathway for obtaining strain fields, with the potential to reconstruct high-resolution maps from individual strain gauge readings. Subject to future validation, the framework could address the limitations of subsea DIC operation [13].

2. Related Work

This section (i) briefly introduces DIC, outlining its principal applications and operational limitations; (ii) reviews machine learning approaches proposed to accelerate DIC, highlighting their reliance on high-resolution images; and (iii) surveys real-world sensor-placement optimisation in structural health monitoring and related fields.

2.1. Introduction to DIC, Possible Applications and Limitations

DIC is a non-contact optical technique that estimates the surface deformation gradient from images of a speckled specimen, computing high-spatial-resolution displacement and strain maps [14]. DIC works by tracking the movement of pixels in camera images, allowing for the interpolation and generation of a complete deformation field. For precise displacement measurement on a specimen's surface, the image is divided into unique pixel subsets. The sum of squared differences (SSD) criteria, such as the zero-mean normalized sum of squared differences (ZNSSD) used in this study, are employed to maximise the cross-correlation of a pixel subset between images. Thanks to the non-contact, full-field coverage attributes of DIC, it has been applied in numerous full-scale tests of composite structures [15].

2.1.1. DIC Applications in Wind and Tidal Energy Sectors

DIC has been widely used in wind turbine blades for characterising the dynamics of structures [16], health monitoring [17], as well as static and fatigue tests [18–20]. A DIC application in the wind sector, presented in [21], shows how DIC has been used for in situ deflection monitoring of a 3.2 MW wind turbine, with a rotor diameter of 114 m. Two cameras are placed at 205 m from the turbine, and the in-plane, out-of-plane and torsional deflections are computed in image post-processing. They are in good agreement with the results of aeroelastic simulations carried out for the wind turbine blade, based on simultaneous wind measurements. Similarly, authors in [22] measure the blade deflection of a 2.5 MW, 80 m diameter wind turbine. Four cameras placed 220 m away from the turbine are used in the on-site experiment. The research findings show that the displacement can be measured with an absolute error of 25 mm. The work done by [17] shows the application of DIC to compute both strain and displacement maps for a 5 kW, 4 m diameter wind turbine. The DIC measurements are subsequently used for real-time, on-site health monitoring. The study by [23] presents the application of DIC to a use case similar to the one at FastBlade, where two cameras are used to monitor the strain distribution on a 14.3 m long wind turbine blade, undergoing cyclic stressing. The paper describes the test setup in detail, including camera spacing and DIC operation settings, providing a valuable reference for tests carried out at FastBlade.

A review of the state-of-the-art structural testing of tidal turbine blades [24] mentions the benefits of using DIC for monitoring the properties of tidal turbine blades under test. It also states that the method has been applied to measure strain and deflection in the test centre in Galway, Ireland [25,26], as well as to measure the strain distribution across the blade in Newcastle, UK [27] and at FastBlade [9].

2.1.2. Relevant DIC Limitations

DIC has proven effective for monitoring large composite structures such as wind and tidal turbine blades. Its practical deployment at scale is constrained, however, by three factors: (i) the computational burden of subset correlation and strain calculation grows with image resolution and field of view; (ii) acquisition of high-resolution images imposes high bandwidth and storage demands, which is problematic for long fatigue tests; and (iii) optical requirements preclude subsea operation. It is worth noting that efforts have been made to pursue underwater DIC, including the development of refraction-compensation algorithms [28], sealed binocular systems [29], and broader feasibility studies [13]. Ocean conditions, however, such as turbulence, biofouling, and variable illumination, remain significant challenges to reliable imaging.

2.2. The Use of Machine Learning in DIC

Machine learning (ML) has been incorporated into DIC pipelines primarily to accelerate the computation of full-field displacement and strain maps from raw images by bypassing the subset correlation step. In [11], a convolutional neural network (CNN) computes displacement fields directly from speckle images, achieving inference speeds two orders of magnitude faster than conventional algorithms. A similar approach has been demonstrated in [12,30], where the authors use a synthetic dataset to train CNNs to reconstruct both displacement and strain fields, corresponding to the given speckle pattern. The authors of [31] suggest that combining physical knowledge with ML can be used to expand on the black-box approach of the existing models for displacement and strain map reconstruction, and report increased prediction accuracy. It is worth highlighting that no studies have been identified which incorporate the use of ML models to reconstruct high-resolution displacement and/or strain fields from sparse measurements.

The Limitations of the Existing ML Approach

ML-based DIC methods have primarily been developed to accelerate the conventional pipeline by mapping full-resolution speckle images directly to displacement/strain fields, bypassing iterative subset correlation. Despite the gains of the referenced methods, several practical challenges remain for large-scale structural testing. They assume that systems are equipped with suitable bandwidth and storage to manage the collection of high-resolution raw camera images. While it might not constitute a problem in certain test arrangements, monitoring larger specimens, such as tidal turbine blades, presents an inevitable challenge if high-resolution images need to be recorded over extended periods. Therefore, the reconstruction of high-resolution fields from sparse physical measurements is identified as a gap in ML-related DIC research. Moreover, networks trained on synthetic data may degrade when provided with actual camera readings, requiring re-training. In this work, we use test data to explore the efficacy of methods capable of reducing the amount of experimental data collected, such as the optimisation algorithms for sensing locations, which are introduced in the following sections.

2.3. Sensor Placement Optimisation

Numerous examples across the natural sciences and engineering demonstrate that optimising sensor locations enables accurate reconstruction of high-resolution fields from sparse measurements, when historical system data, typically a low-dimensional basis extracted from past measurements, is used as a prior. In many studies, the selection step is performed via QR factorisation with column pivoting (QRCP) [32–36]. The rationale for this choice is outlined in Section 3.3.1.

The work presented in [32] introduces the paradigms of sparsity and compressed sensing as a preliminary to tailored sensing, which results in better sensor optimisation for reconstruction thanks to the use of historic measurements. The mathematical formulation of the problem is given, and the computation of optimal sensor locations is presented for the reconstruction of fluid flow fields past a cylinder, photographs of human faces, or ocean surface temperature maps. The reconstruction results are shown and compared against alternative sensor selection methods, proving the superiority of applying QRCP in related problems. [33] further extends this method to cases where optimising the location of sensors is influenced by an associated cost function, showing applications to the same real-life data as [32].

QRCP has been applied to reconstruct high-resolution radio-frequency atomic magnetometer scans of an indented aluminium plate in a non-destructive testing application, supporting the accelerated calibration of the test setup [37]. The publication by [35] presents a similar application to the objective of our work, where the distribution of strain sensors is optimised on the surface of an antenna. The surface strain readings are subsequently transformed into the deformation of the shape. The authors apply an iterative sensor placement technique, where the starting configuration is obtained with the aid of QR factorisation, and the following sensors are added using a sequential sensor placement algorithm.

A high-precision, industrial application is presented in [36], which shows the use of the sensor placement algorithm based on QRCP to predict the locations for shims, inserted to enhance the structural rigidity of a composite Boeing aircraft wing. The locations are determined using high-resolution laser wing scans. To enhance the method's performance, the authors compute robust principal component analysis (RPCA) on the dataset before optimising the sensing locations. RPCA is a method of estimating the dominant principal components with strong insensitivity to outliers [38], and its application yields a training dataset with a reduced noise level, enhancing the detection of low-level data patterns.

The authors set an absolute reconstruction error threshold and reach the reconstruction accuracy of 99%, using around 3% of the original scan locations.

These studies suggest that data-driven placement using QR factorisation is a reliable and efficient route to sparse sensing for high-dimensional fields. The presented work adopts the same philosophy for tidal turbine blades at FastBlade, as a low-rank basis derived from historical DIC maps, denoised via RPCA, supports the optimisation of sparse sensing locations, enabling reconstruction of high-resolution deflection and strain fields from a fraction of original measurements.

3. Method

The literature review establishes that DIC is a promising tool for monitoring tidal turbine blades. To address the problems of high bandwidth and significant computational requirements, we avoid collecting full-resolution images by optimising sensing locations. This section presents the associated pipeline, including data pre-processing and reconstructing the high-spatial-resolution displacement and strain maps from sparse measurements. The reader is encouraged to refer to the corresponding section in the Supplementary Materials, where the RPCA (featuring the alternating directions method (ADM) [39]) and sensing location optimisation derivations are described in more detail.

3.1. Dataset Pre-Processing

A single system measurement, denoted by \mathbf{x} , is a 2D map computed through DIC. Since the same sensing location optimisation method can be applied to both displacement and strain maps, \mathbf{x} can represent either of them. First, all measurements must be pre-processed to eliminate any non-numeric values, referred to as NaNs, which appear in the data in areas where the DIC cannot compute displacement information. Their presence prohibits operations such as QRCP from being performed on the dataset. Therefore, the following operation on \mathbf{x} is defined:

$$f(\mathbf{x}) = \mathbf{x}_f, \mathbf{n} \quad (1)$$

where f transforms the original 2D map into a 1D array \mathbf{x}_f , removing any NaN entries from it, and the array \mathbf{n} contains NaN locations within \mathbf{x} . Once the NaN entries have been removed from the flattened measurements \mathbf{x}_f , they are concatenated as columns to form the dataset \mathbf{X} , used to optimise the sensing locations. \mathbf{X} has the dimensions of $m \times n$, where m is the length of \mathbf{x}_f , and n is the total number of maps recorded. Each flattened measurement \mathbf{x}_f can be subsequently reshaped into a 2D map again through the operation f' :

$$f'(\mathbf{x}_f, \mathbf{n}, a, b) = \mathbf{x} \quad (2)$$

where a and b are the original width and height of \mathbf{x} .

3.2. Robust Principal Component Analysis

The RPCA technique can be used to remove outliers from a dataset of high-dimensional samples, even when some of the dataset entries are obliterated. The published applications involve processing fluid flow fields [40], human faces, or surveillance camera images [38]. RPCA decomposes the matrix containing raw data, \mathbf{X} , into the sum of two other matrices, \mathbf{L} and \mathbf{S} :

$$\mathbf{X} = \mathbf{L} + \mathbf{S} \quad (3)$$

where \mathbf{L} is a low-rank matrix, implying that it is characterised by dominant coherent structures and thus contains highly correlated data, and \mathbf{S} is sparse, isolating the outliers

present in \mathbf{X} . In our research, \mathbf{X} represents the raw dataset of displacement/strain measurements, and \mathbf{L} corresponds to the dataset with removed outliers (for further details see Supplemental Material).

3.3. Sensor Placement Optimisation for Reconstruction

3.3.1. Tailored Sensing

The work by [32] defines tailored sensing as the strategic choice of individual sensing locations in a system to reconstruct a full-resolution image with high accuracy, provided that a sufficient number of historical measurements of the complete system are available. When r optimal sensing locations are considered, the problem is formulated in the following way:

$$\mathbf{y} = \mathbf{C}\boldsymbol{\psi}_r\mathbf{a} \quad (4)$$

where \mathbf{C} is the sensing location matrix, \mathbf{y} contains r corresponding system measurements, $\boldsymbol{\psi}_r$ is a reconstruction basis derived from historical measurements, and \mathbf{a} is the vector used for image reconstruction. An observation made by [32] is that near-optimal sensing locations \mathbf{C} can be found through computing QRCP on $\boldsymbol{\psi}_r$. QRCP is a process of decomposing a matrix into an orthonormal matrix and an upper triangular matrix. The pivot locations are selected in an iterative process, where the entry with the largest L^2 norm is considered [41], and the remaining entries are orthogonalised with respect to it. Using this process, system locations with minimal information redundancy are found, which is desirable for sensing location optimisation.

The basis $\boldsymbol{\psi}_r$ can be constructed from the dataset \mathbf{X} using a dimensionality reduction technique to obtain features representative of the entire dataset. In this work, singular-value decomposition (SVD) is used. SVD decomposes the matrix into a product of three matrices, where \mathbf{U} and \mathbf{V} are orthonormalised matrices of row and column singular vectors, respectively, and $\boldsymbol{\Sigma}$ is a diagonal matrix of singular values [42]. It is also possible to approximate \mathbf{X} using just the first r features:

$$\mathbf{X} \approx \mathbf{U}_r\boldsymbol{\Sigma}_r\mathbf{V}_r^T. \quad (5)$$

The linear combination of the columns of \mathbf{U} can be used to reconstruct any sample from the original dataset \mathbf{X} . Thus, \mathbf{U}_r entries are an approximation of the system and can be adapted on a system-specific basis $\boldsymbol{\psi}_r$.

Considering the Equation (4), once the system measurements \mathbf{y} are collected, the only remaining unknown \mathbf{a} can be obtained from the relation. Successively, the reconstruction of the original high-dimensional measurement, $\hat{\mathbf{x}}_f$, can be computed:

$$\hat{\mathbf{x}}_f = \boldsymbol{\psi}_r\mathbf{a}. \quad (6)$$

For further details, see Supplemental Material.

3.4. Processing Pipeline

The processing pipeline incorporated in the paper is outlined in Figure 2, with the training and reconstruction parts highlighted. The pipeline applies to the reconstruction of both displacement and strain maps computed with DIC (for further details, see Supplemental Material).

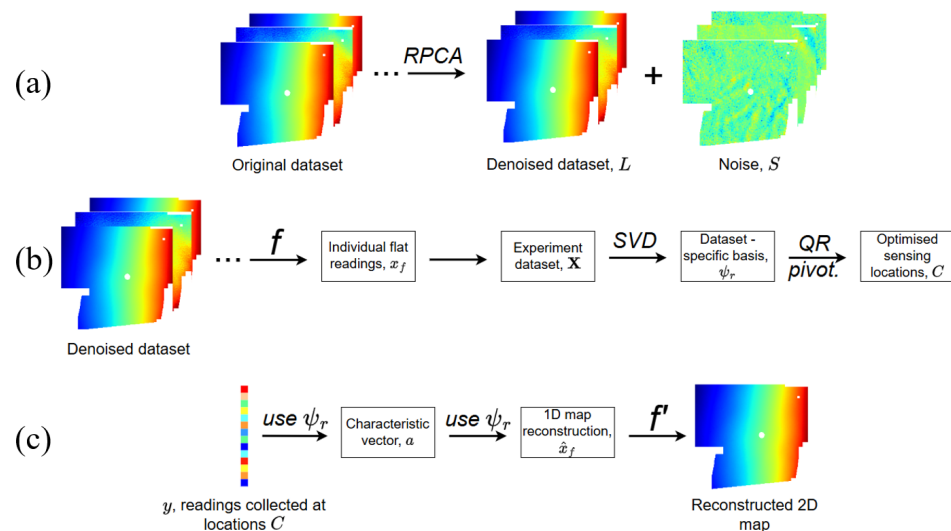


Figure 2. (a) The preprocessing pipeline: RPCA is applied to the original dataset, effectively denoising the maps. (b) The training pipeline: The original dataset (where white areas represent NaN entries) undergoes flattening, the NaN entries are deleted, and the 1D matrices are concatenated. SVD is used to construct a dataset-specific basis ψ_r , which is subsequently used to determine the optimised sensing locations C. (c) The reconstruction pipeline: New system readings are taken, and each set of readings can be used to reconstruct a full-dimensional image using ψ_r , which is then reshaped and populated with NaNs through f' .

4. Results

4.1. The Dataset

This subsection describes the parameters of the experimental setup and the basic properties of the collected dataset.

4.1.1. Experimental Setup

The dataset used in this work was collected during a composite tidal turbine blade test, in which the blade was deflected by three actuators working at a frequency of 1.0 Hz. The maximum and minimum loads exerted by the blades within one cycle were 65 kN and 15 kN, respectively. The blade was 5.25 m long and weighed 1588.59 kg. It also had a natural frequency of approximately 18 Hz, and a cross-section described by the NACA 63-4XX aerofoil series. The thickness-to-chord ratio of the blade varied from 55% at the root to 18% at the tip. The blade’s innermost portion had a cylindrical cross-section with a thickness-to-chord ratio of 100% [10]. The experimental setup is presented in Figure 1b.

During a portion of the fatigue test, two camera pairs were used to record raw images of the speckle pattern in the middle section of the tidal turbine blade between two support saddles. The images were recorded at a frequency of 10.0 Hz, and the corresponding displacement and strain maps were calculated during the post-processing process. The associated DIC parameters are presented in Table 1, and the detailed calibration parameters can be found in Table A1 in Appendix A.

Table 1. DIC parameter settings.

Subset size	27 [pix]
Step size	9 [pix]
Correlation criterion	ZNSSD [-]
Interpolation function	Local Bicubic Splines [-]
Shape function	Quadratic [-]

Table 1. Cont.

Stereo transformation	Quadratic [-]
Pre-filtering	Gaussian
Strain window	29 [-]
Virtual strain gauge	279 [pix]
Strain formulation	Hencky

4.1.2. Virtual Strain Gauge Size Selection

Among the parameters presented in Table 1, the virtual strain gauge (VSG) size deserves particular attention, as its dimensions may have a significant impact on the strain field noise computed by DIC. It is also noted that the VSG size is related to the strain window size by the following relationship:

$$VSG_{size} = (window_{size} - 1) \times step_{size} + subset_{size} \tag{7}$$

where all considered values are presented in Table 1. The VSG size is determined from analysis of the longitudinal strain (ϵ_{xx}), which is dominant for the specimen. Individual pixels in the region of interest (ROI) are considered, and the resulting ϵ_{xx} values are computed for the changing VSG size. In line with the instructions provided in relevant guides, the VSG for which the highest strain value is obtained ought to be chosen. An example of the obtained ϵ_{xx} variation is presented in Figure 3a. While the analysis of individual pixels is likely to yield different VSG sizes due to data variation, the other part of the VSG optimisation process involves noise analysis across a larger portion of the blade. Therefore, a line is drawn along the neutral axis of the specimen, and different VSG sizes are used to compute strain values for each pixel along the line. A sample process is presented in Figure 3b, where the ϵ_{xx} variation is observed to confirm the theoretical prediction, which states that increasing the VSG size contributes to noise reduction. On the other hand, a larger VSG decreases spatial resolution [43]. Considering the data presented in Figure 3a,b, the VSG size of 279 is adopted in the study.

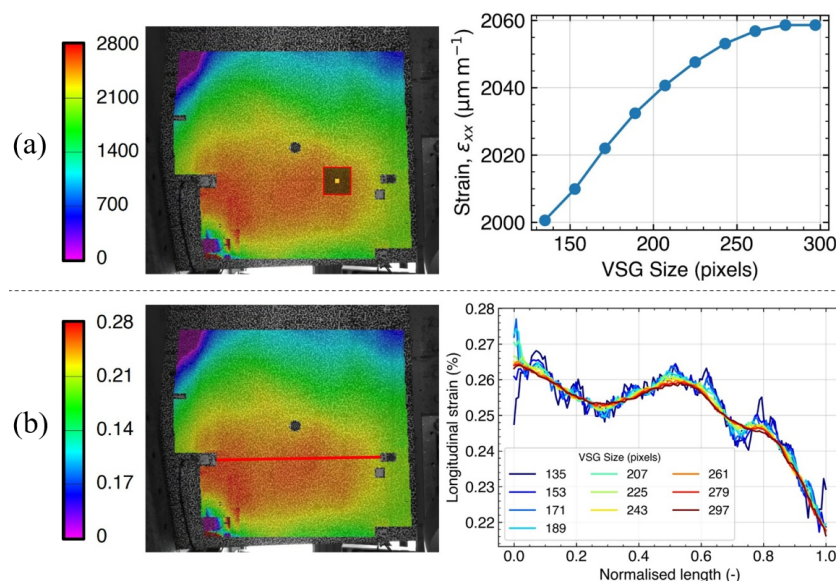


Figure 3. (a) The ϵ_{xx} distribution for the single pixel marked, and the corresponding longitudinal strain variation for a changing VSG size. (b) The ϵ_{xx} distribution with the central axis marked, and a corresponding longitudinal strain variation for ten VSG sizes considered.

4.1.3. The Datasets

The experimental data comprises 1374 sets of camera images, for which individual displacement and strain maps can be computed. The DIC setup is capable of computing displacement magnitudes in three directions, denoted as u , v , and w , as presented in Figure 4. Moreover, the total displacement, d , is computed using the three individual components as:

$$d = \sqrt{u^2 + v^2 + w^2}. \tag{8}$$

In this work, the distribution of d is considered in the sensing location optimisation process, as the information about the total displacement magnitude is more valuable than any of the individual components alone.




Figure 4. The rendered image of the tested blade’s model. The naming convention for the three directions relative to the blade’s surface is provided. The section of the blade investigated in this work is marked in blue. The dashed line marks the central axis of the blade.

Strain values, which are calculated directly from the displacement distribution in three directions, include axial strains ϵ_{uu} and ϵ_{vv} , and the shear strain ϵ_{uv} . These values can be used to calculate the maximum principal strain, ϵ_1 using [44]:

$$\epsilon_1 = \frac{\epsilon_{uu} + \epsilon_{vv}}{2} + \frac{\sqrt{(\epsilon_{uu} - \epsilon_{vv})^2 + \epsilon_{uv}^2}}{2}. \tag{9}$$

The maximum principal strain represents the largest possible normal strain acting in a material [45]. Therefore, it is considered in the optimisation process in this work, as it provides the most conservative means to assess the feasibility of a structure to withstand loads.

The reconstructions for displacement and strain are reported separately in Sections 4.3 and 4.4. To obtain an unbiased assessment, sensing locations are learned strictly from a training set and evaluated on a disjoint validation set. For each dataset (d and ϵ_1), which contains 1374 maps, we randomly reserve 300 maps for validation and use the remaining 1074 for training. Moreover, random permutation (with a fixed seed) is applied to both datasets to break any temporal order in the acquisition sequence.

4.2. Preprocessing the Experimental Data

Quantifying the level of noise in the displacement and strain maps is inherently difficult due to the lack of ground truths. Therefore, the nature of noise and its impact on measurements can, to a large extent, only be speculated upon based on the available information on DIC’s parameters. It is, however, possible to quantify the impact of different denoising techniques on displacement and strain maps through the lens of low-rank

representation. To this end, the original maps are compared with results computed using RPCA (introduced in Section 3.2) and Gaussian filtering, a commonly applied technique to counteract random, low-amplitude noise. The extent to which a given dataset can be explained by low-rank components can be quantified by performing principal component analysis (PCA) and constructing a relevant scree plot, which shows the amount of total variance explained by a growing number of principal components. The scree plots for the displacement and strain data are presented in Figure 5a and 5b, respectively.

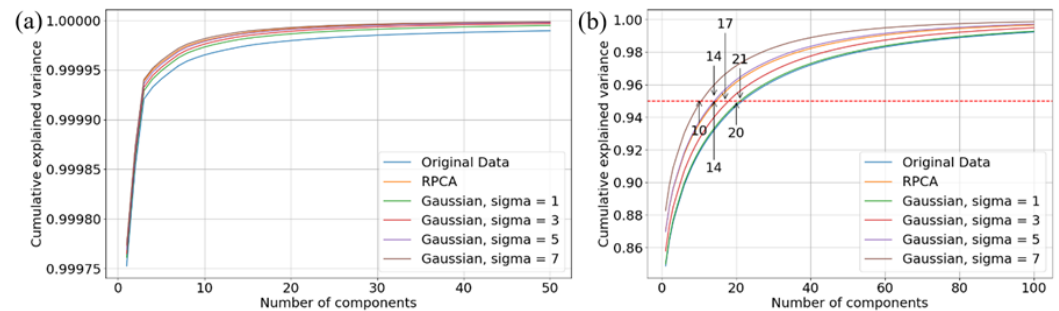


Figure 5. Scree plots computed for the experimental (a) displacement, and (b) strain data. The graph showing the variance of strain includes the 95% threshold and the number of principal components required for each distribution to reach it.

Looking at Figure 5a, it is evident that the variance explained by the principal components computed for the original dataset is smaller than that for the filtered data. RPCA improves the amount of variance explained relative to data subject to Gaussian filtering with σ values of 1, 3, and 5, but explains less variance than when $\sigma = 7$ is used. It is observed, however, that even with a single principal component, the variance explained exceeds 99%, a commonly used threshold in related problems. A more detailed look at the impact of denoising on the data is presented in Figure A1 in the Appendix B, where the difference between the original and denoised values is computed for the entire displacement dataset, and the distribution of this difference is plotted as a histogram. The difference between the original and low-rank maps computed with RPCA is significantly smaller than that with Gaussian filtering, with virtually all values falling into the bin closest to zero. This suggests that, unlike using pre-set filters, RPCA promotes a low-rank representation without blurring.

The analysis of cumulative explained variance for the strain data (Figure 5b) confirms that the variance in the original data is better explained by the principal components computed from the denoised data. In Figure 5b, there is also the number of principal components indicated, which is required to surpass the threshold of 95% of explained variance. RPCA achieves the same score as Gaussian filtering with the σ value of 5, requiring 14 principal components to pass the threshold. The corresponding histogram visualising the distribution of differences computed between the original and denoised images (Figure A2 in the Appendix B) confirms a similar trend to that described for displacement, as the differences recorded for RPCA are, on average, significantly smaller.

4.3. Displacement

Following the discussion in Section 4.2, RPCA has been applied to the displacement data. Figure 6 shows the results of the RPCA process for three randomly chosen samples.




Figure 6. The presentation of the raw d maps, x , and the associated RPCA results, where l is the low-rank component, and s is the resulting sparse matrix. Rows (a–c) represent three randomly chosen samples. The displacement units are millimetres, and each subplot has a corresponding colourbar.

Figure 6 shows that the values of s , comprising patterns that are poorly explained by low-rank components, are small relative to the original displacement values. Moreover, the visual inspection of the de-noised images, l , does not show a significant change in either the magnitude or the distribution of displacement relative to the original maps, x .

4.3.1. Optimising Sensing Locations

The sensing locations are optimised using the method presented in Section 3.3.1. Although the visual inspection of the original and low-rank displacement maps computed with RPCA shows no major differences, the sensing location optimisation is attempted for both X and L . It is worth highlighting that the sparse measurements used for reconstruction (denoted as y in Section 3.3.1) are always extracted from the original maps X , since it is assumed that data captured during normal operation would be noisy. To quantify the reconstruction accuracy, relative reconstruction error, δ_{rel} , is computed using the following relationship:

$$\delta_{rel} = \frac{\text{mean}(|\hat{x}_f - x_f|)}{\text{mean}(|x_f|)}. \tag{10}$$

In the experimental procedure, a range of possible sensing location numbers is investigated. For each of these numbers, reconstructions for all maps in the validation dataset are computed. Once the corresponding reconstruction error has been calculated for all 300 samples, δ_{rel} values are averaged across all the samples to yield the averaged reconstruction error for a certain number of sensing locations. The resulting relationship between the number of sensing locations and the average reconstruction δ_{rel} for the entire validation dataset is presented in Figure 7.

Figure 7 shows that, despite there being seemingly little difference between the raw displacement maps and their low-rank representations, the reconstruction error is significantly different for the X and L datasets. The non-processed dataset yields a higher reconstruction error for all investigated numbers of optimal sensing locations. While the reconstruction error for low-rank images generally decreases with the increasing number of sensing locations, the relationship is significantly less stable for the unprocessed data. For the purpose of more detailed analysis, the points for chosen sensor numbers are marked

in Figure 7. They are referred to as elbow points, and constitute sensor configurations for which the average δ_{rel} reaches a local minimum, prior to its increase over the subsequent sensor number increment. The points with a higher average δ_{rel} relative to the previous elbow point have been left out. The reconstruction error values for these points are shown in Table 2.

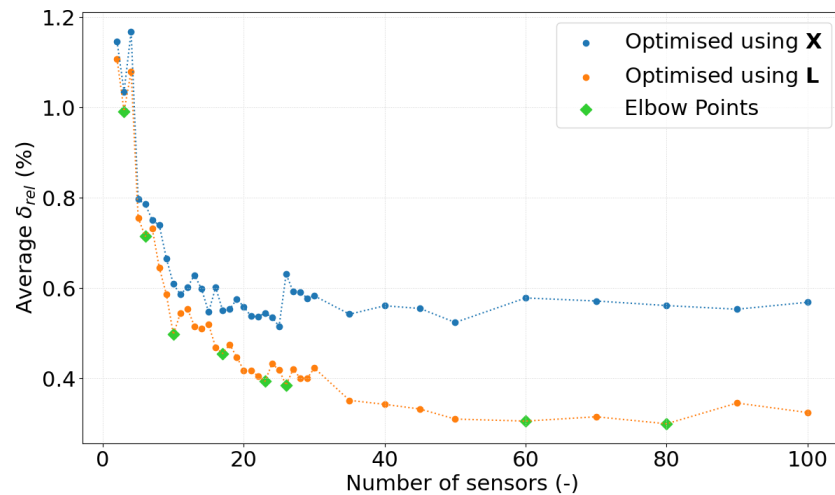


Figure 7. The change in the average δ_{rel} varying with the number of sensing locations, computed for d maps. One series of points was optimised for the dataset of raw images, X , and the other one for the low-rank maps, L . Elbow points marked on the latter series represent sensor numbers for which a significant drop in the reconstruction error is observed.

Table 2. Elbow point values recorded for the d reconstruction error variation.

Number of Sensors	Average δ_{rel}
3	0.99%
6	0.71%
10	0.50%
17	0.46%
23	0.39%
26	0.38%
60	0.31%
80	0.30%

It has to be noted that the DIC measurement uncertainty is difficult to quantify, particularly in an industrial setting [46]. Measurement error can stem from a variety of sources, such as lens distortion, air turbulence, or image noise [47]. Considering the test setup presented in this work, the uncertainty can be further magnified by camera vibration, although care has been taken to minimise this undesired motion by installing the cameras on a separate scaffolding. As a result, the accuracy of the displacement maps computed from δ_{rel} can exceed the achievable DIC measurement accuracy. It is therefore highlighted that the error in this work has been quantified relative to the DIC maps, and not the displacement ground truth. The reconstruction accuracy relative to the actual blade displacement would be a combination of the DIC setup accuracy and the reconstruction accuracy, thereby magnifying δ_{rel} measurement uncertainty relative to the δ_{rel} values demonstrated in this work.

To investigate the reconstruction results more deeply, the actual reconstructions for some of the elbow points are presented in Figure 8. The associated sensor patterns are presented and discussed in Appendix C.1.

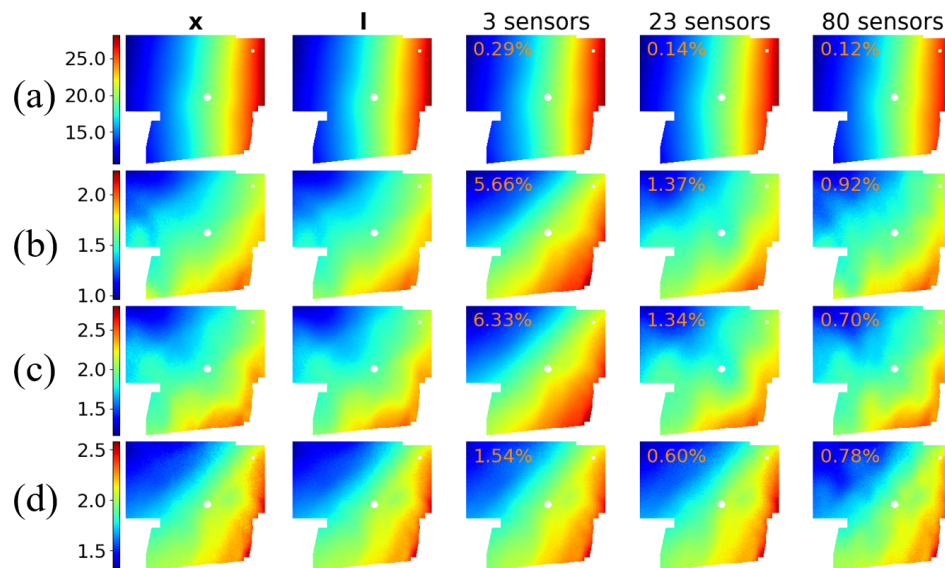


Figure 8. The raw d map, x , the corresponding low-rank component, l , and high-resolution reconstructions computed for 3, 23, and 80 sensing locations. Rows (a–d) represent four samples chosen for illustrative purposes. The reconstructed samples have their relative reconstruction error value plotted in the top-left corner.

Figure 8 shows selected reconstructions for 3, 23, and 80 sensors. The samples in the figure are chosen to demonstrate the following events:

- Case (a): all reconstructions have high accuracy, and the error decreases with a growing number of sensing locations;
- Cases (b) and (c): all reconstructions have relatively high error, but it is still observed to decrease with an increasing number of sensing locations;
- Case (d): Three sensing locations still result in the largest error, but 80 sensors result in greater error than 30 sensors.

The visual inspection of the displacement maps in row (a) shows no considerable difference between the original sample and any of the presented reconstructions. Rows (b) and (c) show a more complex displacement pattern, whose reconstruction visibly fails when just three sensing locations are used. The difference between 23 and 80 sensors is subtle, yet visible. The difference between 23 and 80 sensing locations is also visible in row (d), where the former reconstruction case is visibly more similar to the original displacement map.

It can be observed that the high-accuracy reconstructions shown in row (a) are recorded for larger blade deflections than in the other cases. The displacement values closer to the tip of the blade (the right side of the maps) exceed 25.0 mm, while they do not reach 3.0 mm in any other presented case. Better reconstruction results for higher displacements may be a vital consideration for future applications, as high-load cases are more valuable in assessing the blade’s fatigue.

4.4. Strain

The same steps as those described in Section 4.3 are followed to optimise the sensing locations for the case of the maximum strain (ϵ_1) maps. RPCA results for selected strain map samples are presented in Figure 9.

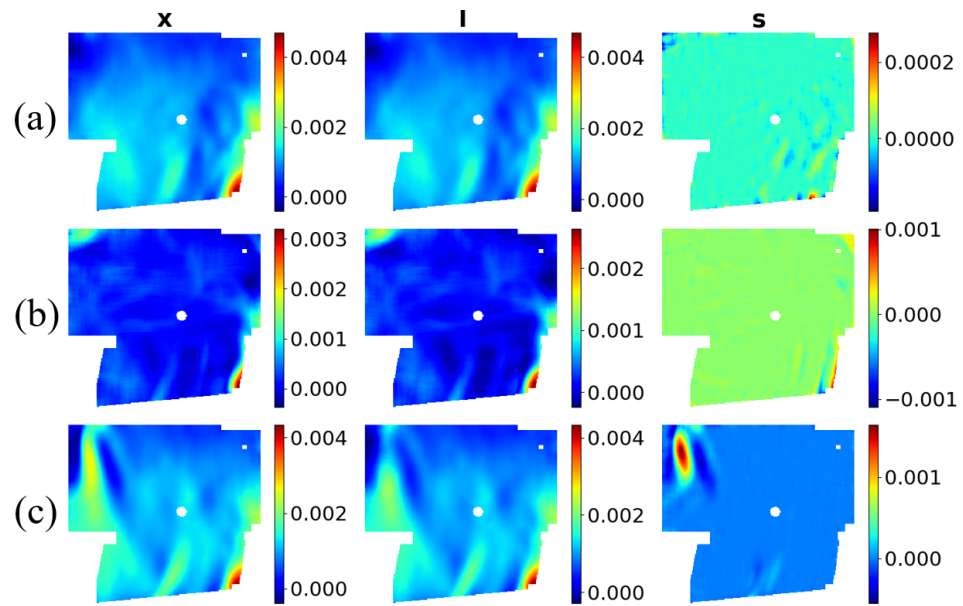


Figure 9. The presentation of the raw maximum principal strain maps, x , and the associated RPCA results, where l is the low-rank component, and s is the resulting sparse matrix. Rows (a–c) represent three randomly chosen samples. Strain values are dimensionless, and each subplot has a corresponding colourbar.

Similar to the displacement maps, the cases of RPCA decomposition presented in Figure 9 show little difference between the original maps and the low-rank maps at visual inspection. The only observable difference is presented in case (c), where higher strain values form a spot in the top-left corner of x , which is removed from the low-rank representation. Subsequently, sensing locations are optimised for chosen r values using the original and low-rank datasets, yielding the results presented in Figure 10.

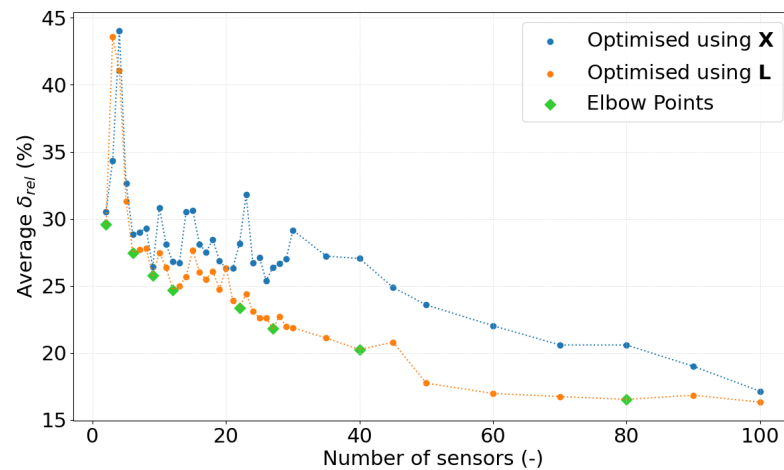


Figure 10. The change in the average δ_{rel} varying with the number of sensing locations, computed for ϵ_1 maps. One series of points was optimised for the dataset of raw images, X , and the other one for the low-rank maps, L . Elbow points marked on the latter series represent sensor numbers for which a significant drop in the reconstruction error is observed.

Similar to displacement maps, elbow points are also marked in the diagram to highlight the step changes in the average δ_{rel} . It is also evident that better reconstruction results have been obtained with the low-rank dataset, L . A major difference between the results obtained for displacement and strain is that the actual average error values are significantly

larger, varying between 43.59% for three sensing locations and 16.31% for 100 sensors. The δ_{rel} values for the elbow points are presented in Table 3.

Table 3. Elbow point values recorded for the ϵ_1 reconstruction error variation.

Number of Sensors	Average δ_{rel}
2	29.56%
6	27.47%
9	25.76%
12	24.70%
22	23.33%
27	21.80%
40	20.24%
80	16.52%

To provide better insight into the properties of the reconstructions, selected strain maps are presented in Figure 11. The associated sensor patterns are presented and discussed in Appendix C.2.

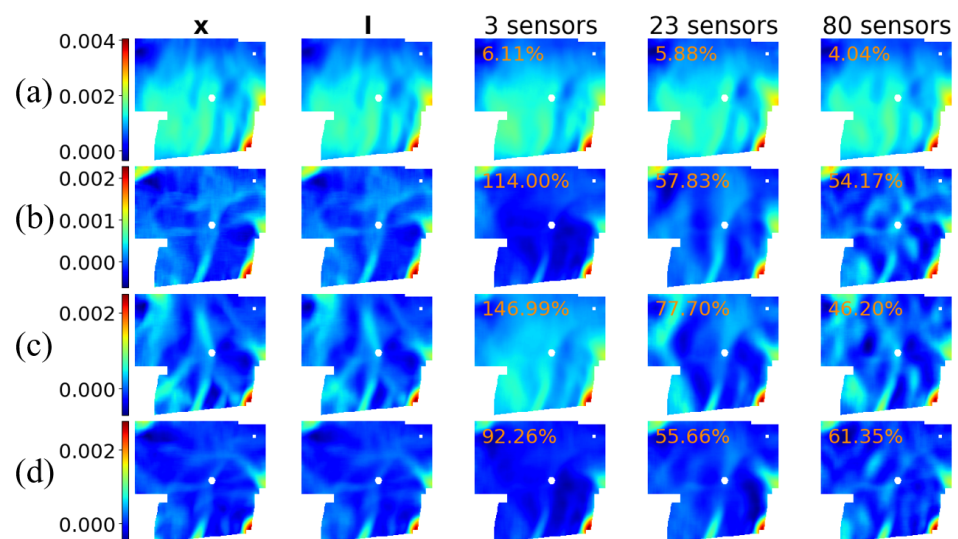


Figure 11. The raw maximum principal strain map, x , the corresponding low-rank component, l , and high-resolution reconstructions computed for 3, 23, and 80 sensing locations. Rows (a–d) represent four samples chosen for illustrative purposes. The reconstructed samples have their relative reconstruction error value plotted in the top-left corner.

Similar to the case of displacement reconstructions, sample (a) presents three decent reconstructions; cases (b) and (c) show some of the poorest reconstructions; and case (d) represents an example where the error does not decrease with an increasing number of sensing locations. The smallest average δ_{rel} is also recorded for values with the largest magnitude, which might have a positive impact on the final application of the method if the inspection of the blade’s condition focuses on the maximum loading case. Moreover, although some of the average δ_{rel} values are very large (exceeding 100%), the general strain distribution patterns are preserved in the reconstructions. Case (d) shows that increasing the number of sensing locations might harm the reconstructions, as it seems that features which are not originally found in the strain map appear in the reconstructed image.

4.5. Real-Time Implementation

One of the major suggested advantages of the method presented is its ability to run in real time, thanks to the collection of data from manually selected, sparse locations on the specimen. The software used for DIC capture in this work is MatchID [48]. The framework presenting the implementation details for a sample specimen is described in the Supplementary Materials attached to this work.

To benchmark the results, and validate the proof-of-concept implemented in real time, the following stages are considered: (a) reading a row of data from the MatchID's output file, representing a single-frame measurement; (b) computing ϵ_1 from individual displacement components (u , v , and w); (c) reconstructing a full-resolution image using the pre-trained reconstruction matrix. The research procedure in this work assumed sensing at most 100 locations per frame, which is why this number was chosen to benchmark the process, presenting the worst-case scenario in terms of computational time.

The aim of the benchmarking is to evaluate the feasibility of the computational method, and not the DIC hardware specifications, such as the cameras, wiring, and the processing computer, as they also constitute a limiting factor in conventional approaches. Therefore, the benchmarking process neglects the time required for physical measurements, image transmission from the cameras, and the output of the individual displacement data points. In the benchmarking test, the strain computation from displacement data was replicated in a Python (release 3.13.11) script, using the same VSG size as that used in the setup described in this work. The benchmarking cycle count was set to 36,000, corresponding to 1 h of running DIC at 10.0 Hz. The benchmarking results, together with the contribution of each stage, are presented in Table 4 below. The results were obtained using a computer at FastBlade with an AMD EPYC 7513 processor and the following specifications: 64 cores, 128 logical processors, and a base speed of 2.60 GHz.

Table 4. Average benchmarking results for 36,000 cycles.

Process	Time (s)	Time Ratio
Loading displacement data	0.09044	95.04%
Computing ϵ_1 values	0.00135	1.42%
Full-resolution field reconstruction	0.00337	3.54%
Total	0.09516	100.00%

The results show that at the current processing rate, the required time is shorter than 0.1 s, which is the maximum permitted time interval for real-time operation at 10.0 Hz. Since the default MatchID setting causes each incoming measurement to appear as a new row in the output file, an important consideration for long tests is the time required to open the resulting data frame. Considering the benchmarking results presented, there is limited time overhead to allow handling larger fields. It is, therefore, advised that ways to overcome this potential overhead are envisaged for the future. These include saving each measurement in a new text file, exploring binary representation to save memory, and parallelising tasks across processing units.

5. Comparison with Other Sensing Location Distributions

In this section, we benchmark the proposed sensing-location optimisation against two practical, unoptimised layouts on the same blade section: (i) random placement, and (ii) uniform grids, mimicking a possible engineer's choice. These baselines test whether the presented sensing location optimisation method improves reconstruction quality for the same sensor numbers.

Moreover, the efficacy of matrix multiplication used in the presented method is tested. To this end, for each sensor layout, full fields are reconstructed with several commonly used methods:

- Matrix multiplication: This is the default approach used with our optimised layout. It projects the sparse measurements onto the dataset-specific basis learned from training data and fills the entire region of interest.
- Nearest-neighbour interpolation: Each pixel is assigned the value of the closest sensed location, with the distances often found using l_1 norm or Manhattan distance [49]. It is fast and can extrapolate, but produces piecewise-constant fields.
- Bicubic interpolation: A standard cubic image-processing scheme that yields smooth fields between measured pixels [50]. However, it cannot extrapolate beyond the convex hull of the sensed points, so any area outside the measurement envelope remains unreconstructed when sensors are sparse.
- Kriging: An interpolation method based on a Gaussian process using prior covariance [51]. In addition to the sparsely sampled data, kriging requires knowledge of a variogram model [52]. The method is valued for its computational efficiency. Two types of kriging have been considered, namely ordinary kriging (OK) and universal kriging (UK). OK assumes that the mean value is locally constant [51]. The UK extends the use of OK by integrating drift parameters, which allow for deterministic fluctuations in the mean [53]. UK is applied in this work to reduce the number of assumptions about the collected data, and regionally linear drift terms are used. The adjustable variogram parameters are sill, range, and nugget, set to 1.5, 115.0 (half of the image's width), and 0.5, respectively. The explanation of the individual parameters can be found in [54].
- Biharmonic inpainting [55]: We evaluated this method for completeness, but it is too slow for real-time use for the number of dimensions considered. Using an 11th Gen Intel® i7 (four physical/eight logical cores), the mean runtime was 8.66 s per frame, compared with 0.29 s for basis projection (about $30\times$ slower). It is therefore excluded from subsequent comparisons.

All layouts and reconstruction methods are evaluated in the same way as in Section 4.3.1. For each sensor number, we reconstruct every map in the validation set and report the average δ_{rel} computed over valid pixels.

5.1. Random Sensor Arrangement

Random sensing locations comprise the same number of sensing locations as the r numbers considered in the earlier sections. As random sensing locations are generated using five random seeds, the lowest of the five obtained reconstruction errors is considered in each case to provide a fair assessment.

5.1.1. Displacement Reconstructions Using a Random Sensor Arrangement

We evaluate random layouts for displacement maps under three reconstruction schemes: matrix multiplication, nearest-neighbour interpolation, and bicubic interpolation. Figure 12 compares these baselines against the previously optimised placement learned from L , showing reconstruction error versus the number of sensing locations for all three reconstruction methods.

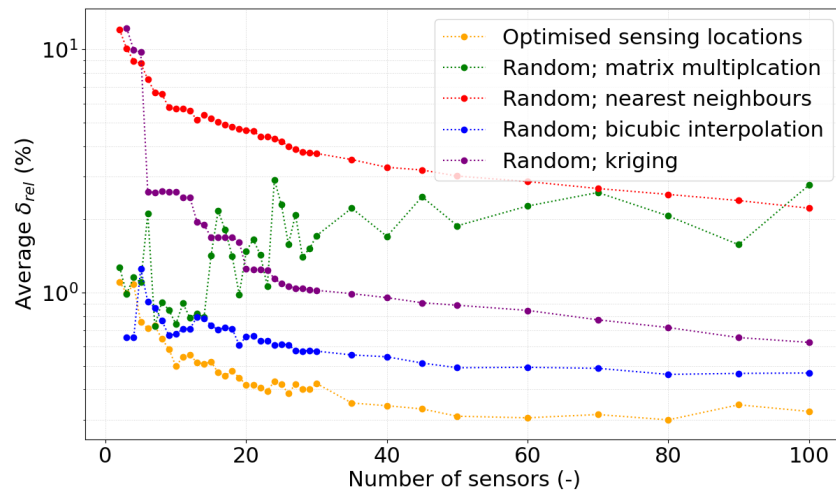


Figure 12. The change in the average δ_{rel} varying with the number of sensing locations, computed for d maps using the optimised sensing locations, as well as random sensing locations, computed with: matrix multiplication; nearest neighbours interpolation; bicubic interpolation; and kriging. To enhance the readability of the plot, the vertical axis is on a logarithmic scale.

The results in Figure 12 show that the choice of reconstruction algorithm has a major impact on reconstruction quality. Apart from the results obtained for random locations using matrix multiplication, most data points show a decreasing average δ_{rel} with an increasing number of sensing locations. Among the random baselines, bicubic interpolation appears the closest, yielding better results than the previously obtained data for three and four sensing locations. However, as previously observed, bicubic interpolation cannot extrapolate beyond the convex hull of the measured points. Consequently, its error is computed only over the subset of pixels that are actually reconstructed, which can make bicubic look deceptively accurate when coverage is sparse. To make this limitation explicit, Figure 13 reports the fraction of reconstructed pixels relative to the full image for bicubic interpolation as the number of sensors increases.

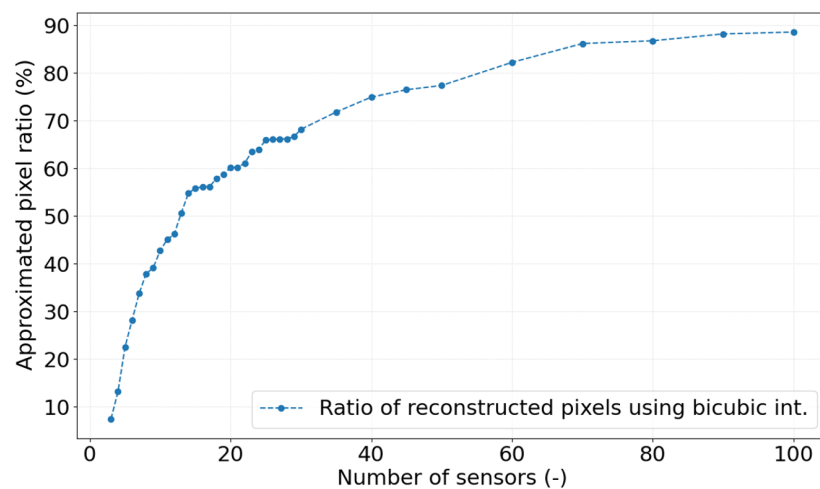


Figure 13. The approximate ratio of the number of pixels reconstructed using random sensing locations and bicubic interpolation to the total number of pixels in the original image, varying with the number of sensing points considered. Each data point is an average ratio computed across all five random seeds considered.

Due to the limitations of bicubic interpolation, the number of reconstructed pixels is low when few sensing locations are used. Beginning at less than 10% for three sensors, 50% of pixels are only reconstructed when 13 sensors are used, and 80% of pixels for

60 sensing locations. In practical terms, despite a low reconstruction error, the results offer limited benefits, as the displacement distribution across only a portion of a structure can be reconstructed.

5.1.2. Strain Reconstructions Using a Random Sensor Arrangement

Using the same method as for displacement (Section 5.1.1), we assess strain-field reconstruction under random sensor placement and compare it with the optimised layouts. The resulting comparison of δ_{rel} , averaged over all 300 samples in the test dataset, is presented in Figure 14.

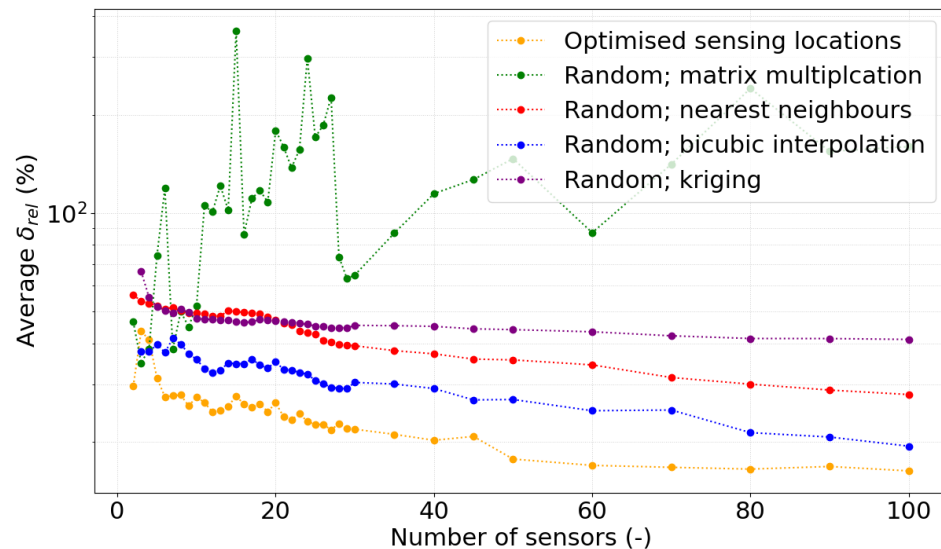


Figure 14. The change in the average δ_{rel} varying with the number of sensing locations, computed for ϵ_1 maps using the optimised sensing locations, as well as random sensing locations, computed with: matrix multiplication; nearest neighbour interpolation; bicubic interpolation; kriging. To enhance the readability of the plot, the vertical axis is on a logarithmic scale.

Except for two very sparse cases ($r = 4$ and $r = 5$), random layouts yield higher errors than the optimised placements across all sensor counts and reconstruction methods. As observed previously (Section 4), strain reconstructions are generally less accurate than displacement reconstructions, reflecting the greater spatial complexity of strain fields and the amplification of noise by numerical differentiation.

Still, the apparent advantage of bicubic interpolation at small r is misleading because it cannot extrapolate outside the convex hull of sensed points. The corresponding coverage curves (fraction of reconstructed pixels) mirror the displacement case in Figure 13.

5.2. Grid Sensor Arrangement

Grids provide even spatial coverage without any prior information about deformation patterns and are straightforward to implement in test campaigns. We evaluate four densities, applied to both displacement and strain experiments: 4×4 , 6×6 , 8×8 , and 10×10 . If a node lies in a masked (NaN) area or outside the region of interest, the corresponding sensor is omitted. As a result, the effective number of sensors may be slightly lower than the nominal grid size. The considered sensor distributions are presented in Figure 15.




Figure 15. Four grid-based sensing location distributions considered in this work. Each subplot has a label comprising the outer dimensions of the grid and the total count of the sensing locations present on its surface.

5.2.1. Displacement Reconstructions Using a Grid Sensor Arrangement

The δ_{rel} values obtained for displacement maps reconstructed with grid-arranged sensors are presented in Figure 16. Each point obtained for the grid arrangement is positioned along the x-axis based on the total number of sensors used in the sensing configuration.




Figure 16. The change in the average δ_{rel} varying with the number of sensing locations, computed for d maps using the optimised sensing locations, and sensing locations arranged in a grid, computed with: matrix multiplication; nearest neighbours interpolation; bicubic interpolation; kriging. To enhance the readability of the plot, the vertical axis is on a logarithmic scale.

First, based on Figure 16, grid layouts paired with matrix multiplication, kriging, or nearest-neighbour interpolation are significantly less accurate than the optimised placements at the same sensor numbers. Second, bicubic interpolation on a grid performs closer to the optimised results, especially as more sensing locations are used. Unlike the random layouts, even spacing improves the convex-hull coverage, so bicubic interpolation reconstructs a larger fraction of the region of interest. Therefore, it is decided to examine the reconstruction error more closely to investigate how it varies for individual test samples.

To compare the methods on a per-sample basis, each grid is paired with a nearby optimised sensor number. When the exact grid count was not used in the original optimisation sweep, we deliberately chose the next lower optimised count to avoid biasing in favour of optimisation, resulting in pairs: 15–15, 30–34, 60–60, and 90–92 sensors. Figure 17 shows the distribution of the average δ_{rel} values across all 300 validation samples for these four pairs.

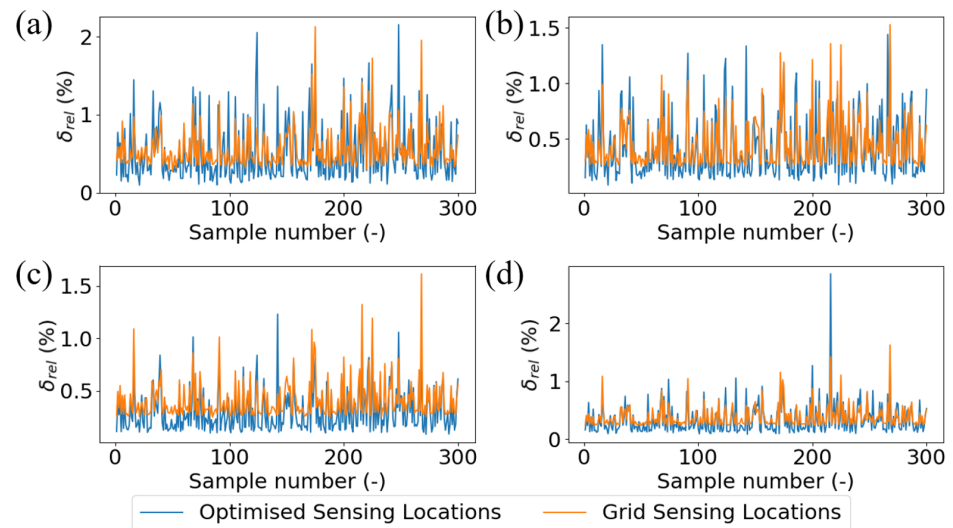


Figure 17. The δ_{rel} variation across 300 samples in the test dataset for optimised sensing locations and grid-arranged sensing locations, with the following sensor counts: (a) 15 sensors each; (b) 30 and 34 sensors, respectively; (c) 60 sensors each; (d) 90 and 92 sensors, respectively.

Across all four comparisons, the optimised placements exhibited a consistent downward offset in error relative to the grid layouts, i.e., lower error for nearly every sample at comparable sensor counts. In addition, matrix multiplication with the optimised layout reconstructs the entire blade section, whereas bicubic calculations on a grid cannot recover the perimeter beyond the measurement envelope. These results indicate that for displacement reconstructions, the proposed optimised placement method provides superior accuracy at the same (or fewer) sensors than uniformly spaced grids.

5.2.2. Strain Reconstructions Using a Grid Sensor Arrangement

We repeat the grid-based analysis for strain fields using the same sensor layouts. The comparison of the average δ_{rel} values computed across all test samples is presented in Figure 18.

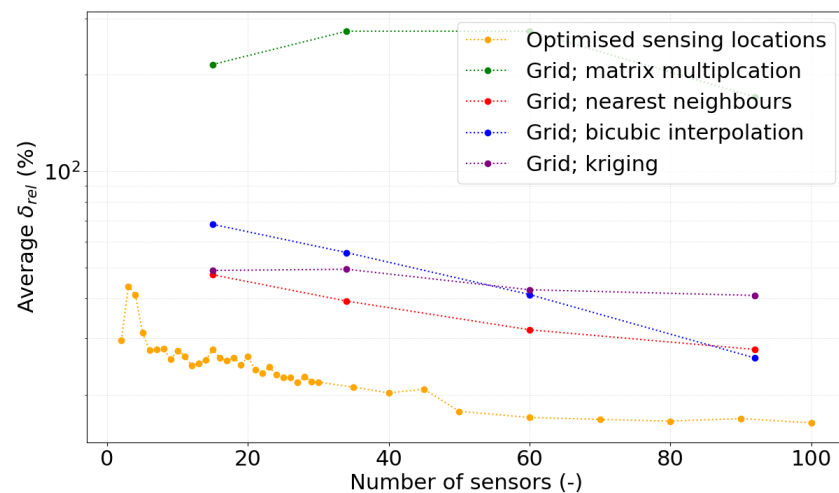


Figure 18. The change in the average δ_{rel} varying with the number of sensing locations, computed for ϵ_1 maps using the optimised sensing locations, and sensing locations arranged in a grid, computed with: matrix multiplication; nearest neighbours interpolation; bicubic interpolation; and kriging. To enhance the readability of the plot, the vertical axis is on a logarithmic scale.

Similar to the case of random sensor arrangement, strain map reconstructions result in greater reconstruction errors than the respective displacement maps. The trace plotted for optimised sensing locations shows that the corresponding reconstruction accuracy exceeds the competing methods, with matrix multiplication resulting in the highest error, and the nearest-neighbour algorithm beating bicubic interpolation and kriging in most cases.

6. Discussion and Future Work

6.1. Time Window and Life Evolution

The sensing locations and reconstruction basis in this study are learned from a finite window of the available fatigue campaign, raising questions about its capacity to capture slow shifts in deformation patterns associated with possible damage accumulation, temperature drift, surface-condition changes, or boundary-condition variations. These limitations are well-recognised in full-scale blade testing [56–58]. To address this concern, we use independent instrumentation data available throughout the extended fatigue campaign (30,000 cycles, equivalent to 20 years of experiencing tide changes), and analyse the relationship between (i) strain-gauge response and (ii) root-bending moment. For more details, see [9].

Using three actuators at fixed blade locations, we observe that the strain-load mapping remains approximately linear considering the peak values after 3×10^4 cycles (see Figure 19). We analysed two strain-gauge locations in the section of the blade covered by the DIC region of interest (ROI): strain gauge 3_1 on the top side and strain gauge 3_4 on the bottom side. For the monitored interval, location 3_1 has a representative goodness-of-fit of $R^2 \approx 0.73$, and location 3_4: $R^2 \approx 0.72$. This supports the assumption that the dominant load–strain relationship is sufficiently stationary for sparse reconstruction over the tested range and is time-invariant under elastic behaviour. Moreover, the reconstruction method presented in this work (see Equation (6)) ensures that the linear scaling of the measured quantities also linearly scales the reconstructed outputs.




Figure 19. Root bending moment vs. strain at the peak of each cycle for locations (a) 3_1 and (b) 3_4.

6.2. Strain Noise Origin

The experimental results show that displacement map reconstruction achieves far better accuracy than the equivalent strain maps. The most obvious reason for this is the increased complexity of the strain patterns relative to the displacement distribution. However, it is possible that strain variations not only represent more complex physical behaviour, but are also significantly impacted by noise. Strain noise originates primarily from displacement uncertainty and differentiation. To find a plausible explanation, it is considered that strain is obtained from displacement measured along the three principal axes, and according to [44], the maximum principal strain ϵ_1 is calculated with the following formula:

$$\epsilon_1 = \frac{\epsilon_x + \epsilon_y}{2} + \frac{\sqrt{(\epsilon_x - \epsilon_y)^2 + \epsilon_{xy}^2}}{2} \quad (11)$$

where ϵ_x and ϵ_y are linear strains in the x- and y-directions, and ϵ_{xy} is the shear strain in the xy plane. Each strain component needed to compute ϵ_1 can be calculated as follows [59]:

$$\epsilon_x = \frac{1}{2} \left(2 \frac{\partial u}{\partial x} + \left(\frac{\partial u}{\partial x} \right)^2 + \left(\frac{\partial v}{\partial x} \right)^2 \right) \quad (12)$$

$$\epsilon_y = \frac{1}{2} \left(2 \frac{\partial v}{\partial y} + \left(\frac{\partial u}{\partial y} \right)^2 + \left(\frac{\partial v}{\partial y} \right)^2 \right) \quad (13)$$

$$\epsilon_{xy} = \frac{1}{2} \left(\frac{\partial u}{\partial y} + \frac{\partial v}{\partial x} + \frac{\partial u \partial u}{\partial x \partial y} + \frac{\partial v \partial v}{\partial x \partial y} \right) \quad (14)$$

where u and v are displacement components. Thus, strain depends on first derivatives and products of derivatives of the displacement field, which could yield higher noise for strain reconstructions than for displacement. On the other hand, since small strains are considered, the product terms are likely to be small and have little effect on the result.

Therefore, it is more likely that the increased noise levels in the strain measurements result directly from the DIC computation process. In the process of matching pixels with varying lighting conditions, among other filtering techniques, the DIC algorithm reduces the displacement noise level by applying a local smoothing function [60]. Additionally, the computation of strain maps from displacement data greatly depends on the size of the virtual strain gauge (VSG), which is a spatial filter for strain computation. Larger VSG sizes result in smoother strain maps, but they introduce bias error and decrease spatial resolution [43]. Therefore, the VSG can be downsized to recover more spatial information, simultaneously raising the risk of noise. The VSG size is thought to contribute the most to the high noise levels observed in the strain distributions considered in this work.

Although RPCA is effective at suppressing outliers in image-like datasets, it is not specifically designed to attenuate the high-frequency measurement fluctuations that become prominent in strain fields. In DIC, strain is obtained through spatial differentiation of the displacement field, which acts as a high-pass operation: both random subpixel uncertainty and systematic interpolation bias in displacement estimates are amplified when computing strain components [61–64]. Consequently, a denoising step that preserves low-rank structure (such as RPCA) can yield only a modest visual difference for strain maps, even when it improves displacement fields.

The fact that the strain reconstructions obtained with the optimal method presented in this work greatly outperform the competing methods is also attributed to the increased complexity of the maps. This is in contrast to the case of displacement, where, in some cases, the grid-like sensor placement can also produce promising results. However, all the comparisons run show that the method presented in this work is superior, which is attributed to the a priori knowledge of the system.

6.3. Possible Extensions to Real-Time Tests

The presented results are promising when up to 100 individual measurement points are considered. Using MatchID, it is possible to manually choose the locations of the recorded pixels. Operating at up to 100 locations, it is thought that all arrangements presented in this work could be feasibly set in the software, providing displacement and strain information alike. Another possible application of the optimisation technique shown in this work is the manual application of physical strain gauges at the optimised locations. Although hardware costs, wiring, and data transmission would cap the sensor count (and thus increase reconstruction error), the ability to reconstruct a high-resolution strain map in service would be unprecedented. We also observe that errors tend to decrease at higher load levels, suggesting that reconstructions under peak operating loads may be especially informative. While further investigation into the direct strain measurement is encouraged, direct in-ocean displacement sensing remains challenging with current technology.

6.4. Limitations of the Proposed Solution and Future Work

All reconstruction errors reported in this work are computed with respect to the DIC-estimated fields, which themselves carry measurement uncertainty. So, a method that favours a low-rank representation can smooth out parts of the data that are actually noise or artefacts from correlation or interpolation, and therefore may yield a small discrepancy relative to the DIC output without implying physically better-than-truth accuracy [62,65]. Accordingly, we interpret small relative errors as consistency with a denoised representation of the DIC-measured field rather than as an absolute accuracy claim. Independent displacement measurements (e.g., LVDT/laser tracking) would provide a valuable external reference and are identified as future validation when instrumentation permits.

The principal limitation of our study lies in its representativeness of the life of the blade, which is constrained by the size and temporal coverage of the DIC dataset used. Section 6.1 aims to justify the credibility of the study, given that the blade continues to operate in the elastic region. However, since both RPCA and the sensing-location optimisation benefit from rich historical data, a limited training set may reduce reliability and may not capture degradation-induced changes in deformation patterns. Therefore, a useful extension to this work would be repeating the same test procedure on data collected throughout the entire fatigue test, so that samples representative of different wear levels are used in the process, and the assumptions made are validated.

Another limitation of the presented work is the fact that the quality of the achieved strain reconstructions might be considered unsatisfactory, as the average δ_{rel} is still above 15% for 100 sensing locations. This may have significant implications due to the presence of strict testing standards (such as DNVGL-ST-0164 [66] for tidal turbines), making this a method which, in its current state, cannot be used for certification purposes. Although there may be reasons to accept lower reconstruction accuracy, given the drastic reduction in recorded data and the theoretical possibility of maximising the amount of strain information collected from strain gauges (such as during underwater operation), it is advisable to explore possible improvements. Apart from incorporating ML (described in Appendix D), it is advised that strain maps are not reconstructed directly, but through the primary reconstruction of displacement fields. In such a scenario, displacement along the three axes can be reconstructed with high fidelity, which can later be processed and used to compute a corresponding strain map, possibly less affected by noise.

7. Conclusions

The major bottleneck for DIC in full-scale blade testing is the acquisition, transmission, and storage of high-resolution images required for displacement and strain estimation. Exploiting redundancy in historical DIC fields makes tailored sensing viable by learning a low-rank basis and optimising a small set of point measurement locations to reconstruct full-resolution maps from sparse data. In a laboratory DIC setup, these points correspond to individually recorded pixels, whose values are used to recover the original-resolution total displacement or maximum principal strain field.

The reconstruction results computed for the dataset of displacement maps show that high-resolution images can be reconstructed with an average δ_{rel} below 1% when at least five sensing locations are used. Due to the way in which strain maps are computed from displacement information, noise levels in the strain maps are magnified, and the resulting patterns are more complex. Therefore, the method reconstructs the strain maps with poorer accuracy, and according to the experimental results, more than 40 sensors need to be used for the average δ_{rel} to decrease below 20%. One of the implications of the high strain reconstruction error is its likely incompatibility with testing standards/protocols, supporting the search for possible improvements of the method.

Comparisons against placement and interpolation baselines show that neither alternative layouts (random or uniform grids) nor alternative reconstruction techniques (kriging, nearest-neighbour or bicubic interpolation) outperform the proposed pipeline at matched sensor numbers. The results closest to the optimised findings were computed for the displacement maps using bicubic interpolation for grid-arranged sensors. However, a closer examination of the error distribution reveals that the proposed method consistently outperforms the competitor for nearly all samples in the validation dataset.

In addition to optimising the location of collected pixels in a fatigue test of a tidal turbine blade, it is suggested that strain map reconstruction could be attempted from physical strain gauge readings. Considering the problem domain, this approach might be applicable to increasing the amount of data collected from a specimen operating underwater. Although the relatively high strain error values are certainly a disadvantage of the method, the ability to obtain high-resolution strain data from in situ blade operation would be unprecedented, and could justify further development and experimental validation. Also, the measured error values tend to be lower for the high-loading cases, which are more important from the maintenance standpoint.

A significant limitation of the presented work is its reliance on data collected during a portion of a fatigue test. It is therefore suggested that data from multiple test periods is incorporated in the training procedure in future work, to account for the changing distributions of displacement and strain. Finally, in light of remarkably better reconstruction performance for displacement maps, it is suggested that instead of optimising sensing locations for strain maps directly, their computation from reconstructed displacement maps is evaluated.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/jmse14050408/s1>, Figure S1: The sample used in the mock live-tracking procedure.; Figure S2: The sample during the calibration procedure, performed using the purpose-designed plate.; Figure S3: A screenshot of the software, showing the manual selection of four points on the surface of the specimen.; Figure S4: A screenshot of the software, showing: (a) image from the first camera with the considered points marked; (b) image from the second camera; (c) real-time total displacement chart for each point considered; (d) total displacement values displayed for each point.

Author Contributions: Conceptualization, M.J.M., M.A.V.C., F.C., C.M.Ó.B. and S.L.D.; methodology, F.C. and S.L.D.; software, M.J.M. and S.L.D.; validation, M.J.M. and S.L.D.; formal analysis, M.J.M., M.A.V.C. and S.L.D.; investigation, M.J.M., M.A.V.C. and S.L.D.; resources, F.C.; data curation, M.J.M., M.A.V.C. and S.L.D.; writing—original draft preparation, M.J.M., M.A.V.C. and S.L.D.; writing—review and editing, F.C. and C.M.Ó.B.; visualization, M.J.M. and S.L.D.; supervision, C.M.Ó.B. and S.L.D.; project administration, S.L.D.; funding acquisition, C.M.Ó.B. and S.L.D. All authors have read and agreed to the published version of the manuscript.

Funding: All authors: This research was supported by the Engineering and Physical Sciences Research Council under the CoTide Programme grant number EP/X03903X/1. M.J.M.: The doctoral research project has been funded by Babcock International and The Data Lab. M.A.V.C.: This research was also supported by the Wind & Marine Energy Systems & Structures CDT under the Engineering and Physical Sciences Research Council grant agreement EP/S023801/1.

Data Availability Statement: The data that support the findings of this study are available from the corresponding author upon reasonable request.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. DIC Calibration Parameters

The retrieved calibration parameters for the DIC setup are shown in Table A1.

Table A1. DIC calibration parameters.

Parameter	Value	Parameter	Value
Speckle size	5.0	No. of distortions	5
Camera 0 Intrinsic Parameters			
Fx	3585.93	Fs	1.96
Fy	3584.25	K1	−0.145
K2	0.188	K3	0.345
P1	−0.0000107	P2	−0.00304
Cx	2008.68	Cy	1041.48
Xoffset	0	Yoffset	0
Camera 1 Intrinsic Parameters			
Fx	3578.09	Fs	−0.191
Fy	3577.43	K1	−0.140
K2	0.0714	K3	1.664
P1	0.000198	P2	0.00431
Cx	2125.04	Cy	1089.60
Xoffset	0	Yoffset	0
Extrinsic Stereo Parameters			
Tx	−553.34	Theta	0.404
Ty	−3.83	Phi	18.45
Tz	60.36	Psi	0.302

Appendix B. Histograms Computed for Denoised Displacement and Strain Data

Figure A1 shows the differences computed between the collected displacement data and the denoised datasets. The low-rank representation computed with RPCA is denoted as L , and the matrices denoted as M contained data which underwent Gaussian filtering, with their subscripts containing the σ value used for denoising.




Figure A1. The histogram of the differences computed between the displacement data and the denoised data, distributed over 50 bins.

Figure A2 shows the differences computed between the collected strain data and the denoised datasets, following the same naming convention as the displacement case.




Figure A2. The histogram of the differences computed between the strain data and the denoised data, distributed over 50 bins.

Appendix C. Sensing Location Patterns

Appendix C.1. Displacement

Figure A3 presents the optimal sensing patterns associated with different numbers of sensors considered. The minimum number of sensors presented is three, and the maximum number is 80. The intermediate cases of 10 and 23 sensing locations are also presented, as for these patterns, it should still be feasible to manually choose pixels to be measured with DIC or any other sensing technique.

All sensing configurations presented in Figure A3 exhibit a wide spread across the considered blade section. When just three sensing locations are optimised, they are placed at the edges of the map. However, when more sensors are used, they appear to be placed along horizontal lines, capturing the biggest changes in the displacement of the blade.




Figure A3. The visualisation of the optimised displacement sensing locations for 3, 10, 23, and 80 sensors. Each sensing location is plotted as a point over the averaged map of all low-rank maps, **L**.

Appendix C.2. Strain

The location of the optimised sensing locations for the chosen elbow points is presented in Figure A4. It can be seen that, in all cases, the sensors are quite evenly distributed across the surface of the blade. However, since the distribution pattern seems to be more complex than in the case of displacement maps, it is harder to assess whether the optimal sensing locations correspond with areas characterised by the largest strain gradients.




Figure A4. The visualisation of the optimised strain sensing locations for 2, 12, 27, and 60 sensors. Each sensing location is plotted as a point over the averaged map of all low-rank maps, **L**.

Appendix D. Machine Learning-Based Enhancement

A further point worth mentioning in the context of the presented research is the effort made to enhance the quality of the reconstructions using ML. The reason for attempting this is not only the ubiquity of ML applications in signal processing, but also a chance for improvement in the area, where ML has already proven to be an effective tool. In the experimental procedure, a ground truth (i.e., the original displacement/strain map) exists for each reconstructed image. In ML terms, translating one image into another image is referred to as image-to-image (I2I) translation [67]. Notably, some of the most effective tools for I2I tasks are pixel-to-pixel (pix2pix) models [68,69]. These models are CNNs, based on a U-net architecture [70], making them suitable for processing tasks where both the input and the output are images.

Strain maps have been used to train a pix2pix model, as they suffer from poorer reconstruction accuracy than the displacement maps. The training dataset of 1074 maps was used to train the model using pairs of original and reconstructed images. The 300 remaining images in the validation dataset were used to quantify the ML model's performance. In addition, the strain values in the images were manually clipped to ensure that no outliers were present in the data, and several ways of filling in the NaN values were considered so that full 2D images could be used for training.

The validation of models has shown that although their application might contribute to a smaller reconstruction error in selected cases, the error has generally increased for the validation dataset. This result suggests that the model is unable to learn the complex strain patterns, likely due to (i) limited training data, and (ii) the relative complexity of strain patterns. The findings suggest that it would be challenging to reconstruct high-resolution fields computed with DIC using sparse physical measurements, possibly explaining why the gap in the ML-based DIC research exists. It has to be stressed, however, that it might be possible to achieve promising results with ML if more experimental data was available.

References

1. Yolcan, O.O. World energy outlook and state of renewable energy: 10-Year evaluation. *Innov. Green Dev.* **2023**, *2*, 100070. [[CrossRef](#)]
2. Osman, P.; Hayward, J.A.; Penesis, I.; Marsh, P.; Hemer, M.A.; Griffin, D.; Sayeef, S.; Nader, J.R.; Cossu, R.; Grinham, A.; et al. Dispatchability, Energy Security, and Reduced Capital Cost in Tidal-Wind and Tidal-Solar Energy Farms. *Energies* **2021**, *14*, 8504. [[CrossRef](#)]
3. Chowdhury, M.S.; Rahman, K.S.; Selvanathan, V.; Nuthammachot, N.; Suklueng, M.; Mostafaeipour, A.; Habib, A.; Akhtaruzza-man, M.; Amin, N.; Techato, K. Current trends and prospects of tidal energy technology. *Environ. Dev. Sustain.* **2021**, *23*, 8179–8194. [[CrossRef](#)]
4. Collombet, R.; Parsons, A.; Gruet, R. *Key Trends and Statistics 2021*; Ocean Energy Europe (OEE): Brussels, Belgium, 2022.
5. Coles, D.; Angeloudis, A.; Greaves, D.; Hastie, G.; Lewis, M.; Mackie, L.; McNaughton, J.; Miles, J.; Neill, S.; Piggott, M.; et al. A review of the UK and British Channel Islands practical tidal stream energy resource. *Proc. R. Soc. A Math. Phys. Eng. Sci.* **2021**, *477*, 20210469. [[CrossRef](#)]
6. Frost, C. *Cost Reduction Pathway of Tidal Stream Energy in the UK and France*; Technical report; Tidal Stream Industry Energiser (TIGER): UK; France, 2022.
7. Melcher, D.; Bätge, M.; Neßlinger, S. A novel rotor blade fatigue test setup with elliptical biaxial resonant excitation. *Wind Energy Sci.* **2020**, *5*, 675–684. [[CrossRef](#)]
8. Eder, M.; Belloni, F.; Tesauero, A.; Hanis, T. A multi-frequency fatigue testing method for wind turbine rotor blades. *J. Sound Vib.* **2017**, *388*, 123–140. [[CrossRef](#)]
9. Lopez Dubon, S.; Cuthill, F.; Vogel, C.; Ó Brádaigh, C.; McCarthy, E.D. A full-scale composite tidal blade fatigue test using single and multiple actuators. *Compos. Part A Appl. Sci. Manuf.* **2024**, *181*, 108140. [[CrossRef](#)]
10. Lopez Dubon, S.; Vogel, C.; Garcia Cava, D.; Cuthill, F.; McCarthy, E.D.; Ó Brádaigh, C.M. A Full-Scale Tidal Blade Fatigue Test using the FastBlade Facility. *Renew. Energy* **2024**, *228*, 120653. [[CrossRef](#)]
11. Duan, X.; Xu, H.; Dong, R.; Lin, F.; Huang, J. Digital image correlation based on convolutional neural networks. *Opt. Lasers Eng.* **2023**, *160*, 107234. [[CrossRef](#)]
12. Yang, R.; Li, Y.; Zeng, D.; Guo, P. Deep DIC: Deep learning-based digital image correlation for end-to-end displacement and strain measurement. *J. Mater. Process. Technol.* **2022**, *302*, 117474. [[CrossRef](#)]
13. Liu, B.; Lan, S.; Li, J.; Fang, Q.; Ren, Y.; He, W.; Xie, H. Digital image correlation in extreme conditions. *Thin-Walled Struct.* **2024**, *205*, 112589. [[CrossRef](#)]
14. Valdivia-Camacho, M.A.; Lopez Dubon, S.; Cuthill, F.; Munko, M.J.; McCarthy, E.D.; Alam, P.; Ó Brádaigh, C.M. Clamping parameters in full-scale tidal turbine blade tests: A case study. *Ocean Eng.* **2025**, *327*, 120722. [[CrossRef](#)]
15. Janeliukstis, R.; Chen, X. Review of digital image correlation application to large-scale composite structure testing. *Compos. Struct.* **2021**, *271*, 114143. [[CrossRef](#)]
16. Baqersad, J.; Carr, J.; Lundstrom, T.; Niezrecki, C.; Avitabile, P.; Slattery, M. Dynamic characteristics of a wind turbine blade using 3D digital image correlation. In Proceedings of the Health Monitoring of Structural and Biological Systems 2012, San Diego, CA, USA, 12–15 March 2012; Volume 8348, pp. 711–719.
17. Wu, R.; Zhang, D.; Yu, Q.; Jiang, Y.; Arola, D. Health monitoring of wind turbine blades in operation using three-dimensional digital image correlation. *Mech. Syst. Signal Process.* **2019**, *130*, 470–483. [[CrossRef](#)]
18. LeBlanc, B.; Niezrecki, C.; Avitabile, P.; Chen, J.; Sherwood, J. Damage detection and full surface characterization of a wind turbine blade using three-dimensional digital image correlation. *Struct. Health Monit.* **2013**, *12*, 430–439. [[CrossRef](#)]
19. Carr, J.; Baqersad, J.; Niezrecki, C.; Avitabile, P.; Slattery, M. Dynamic stress–strain on turbine blade using digital image correlation techniques part 1: Static load and calibration. In *Topics in Experimental Dynamics Substructuring and Wind Turbine Dynamics, Volume 2: Proceedings of the 30th IMAC, A Conference on Structural Dynamics*; Springer: New York, NY, USA, 2012; pp. 215–220.
20. Poozesh, P.; Baqersad, J.; Niezrecki, C.; Harvey, E.; Yarala, R. Full field inspection of a utility scale wind turbine blade using digital image correlation. In Proceedings of the CAMX, Orlando, FL, USA, 13–16 October 2014; Volume 10, pp. 2891–2960.
21. Winthroth, J.; Schoen, L.; Ernst, B.; Seume, J.R. Wind turbine rotor blade monitoring using digital image correlation: a comparison to aeroelastic simulations of a multi-megawatt wind turbine. *J. Phys. Conf. Ser.* **2014**, *524*, 012064. [[CrossRef](#)]
22. Ozbek, M.; Rixen, D.J.; Erne, O.; Sanow, G. Feasibility of monitoring large wind turbines using photogrammetry. *Energy* **2010**, *35*, 4802–4811. [[CrossRef](#)]
23. Chen, X.; Semenov, S.; McGugan, M.; Hjelm Madsen, S.; Cem Yeniceli, S.; Berring, P.; Branner, K. Fatigue testing of a 14.3 m composite blade embedded with artificial defects—Damage growth and structural health monitoring. *Compos. Part A Appl. Sci. Manuf.* **2021**, *140*, 106189. [[CrossRef](#)]
24. Munaweera Thanthirige, T.R.; Goggins, J.; Flanagan, M.; Finnegan, W. A State-of-the-Art Review of Structural Testing of Tidal Turbine Blades. *Energies* **2023**, *16*, 4061. [[CrossRef](#)]

25. de la Torre, O.; Moore, D.; Gavigan, D.; Goggins, J. Accelerated life testing study of a novel tidal turbine blade attachment. *Int. J. Fatigue* **2018**, *114*, 226–237. [[CrossRef](#)]
26. Finnegan, W.; Jiang, Y.; Meier, P.; Hung, L.C.; Fagan, E.; Wallace, F.; Glennon, C.; Flanagan, M.; Flanagan, T.; Goggins, J. Numerical modelling, manufacture and structural testing of a full-scale 1 MW tidal turbine blade. *Ocean Eng.* **2022**, *266*, 112717. [[CrossRef](#)]
27. Gonabadi, H.; Oila, A.; Yadav, A.; Bull, S. Structural performance of composite tidal turbine blades. *Compos. Struct.* **2021**, *278*, 114679. [[CrossRef](#)]
28. Su, Z.; Pan, J.; Zhang, S.; Wu, S.; Yu, Q.; Zhang, D. Characterizing dynamic deformation of marine propeller blades with stroboscopic stereo digital image correlation. *Mech. Syst. Signal Process.* **2022**, *162*, 108072. [[CrossRef](#)]
29. Wu, T.; Hou, S.; Sun, W.; Shi, J.; Yang, F.; Zhang, J.; Wu, G.; He, X. Visual measurement method for three-dimensional shape of underwater bridge piers considering multirefraction correction. *Autom. Constr.* **2023**, *146*, 104706. [[CrossRef](#)]
30. Boukhtache, S.; Abdelouahab, K.; Berry, F.; Blaysat, B.; Grédiac, M.; Sur, F. When Deep Learning Meets Digital Image Correlation. *Opt. Lasers Eng.* **2021**, *136*, 106308. [[CrossRef](#)]
31. Cheng, X.; Zhou, S.; Xing, T.; Zhu, Y.; Ma, S. Solving digital image correlation with neural networks constrained by strain-displacement relations. *Opt. Express* **2023**, *31*, 3865. [[CrossRef](#)]
32. Manohar, K.; Brunton, B.W.; Kutz, J.N.; Brunton, S.L. Data-Driven Sparse Sensor Placement for Reconstruction: Demonstrating the Benefits of Exploiting Known Patterns. *IEEE Control Syst.* **2018**, *38*, 63–86. [[CrossRef](#)]
33. Clark, E.; Askham, T.; Brunton, S.L.; Nathan Kutz, J. Greedy Sensor Placement with Cost Constraints. *IEEE Sens. J.* **2019**, *19*, 2642–2656. [[CrossRef](#)]
34. Cao, J.; Bu, F.; Wang, J.; Bao, C.; Chen, W.; Dai, K. Reconstruction of full-field dynamic responses for large-scale structures using optimal sensor placement. *J. Sound Vib.* **2023**, *554*, 117693. [[CrossRef](#)]
35. Zhou, J.; Cai, Z.; Zhao, P.; Tang, B. Efficient Sensor Placement Optimization for Shape Deformation Sensing of Antenna Structures with Fiber Bragg Grating Strain Sensors. *Sensors* **2018**, *18*, 2481. [[CrossRef](#)]
36. Manohar, K.; Hogan, T.; Buttrick, J.; Banerjee, A.G.; Kutz, J.N.; Brunton, S.L. Predicting shim gaps in aircraft assembly with machine learning and sparse sensing. *J. Manuf. Syst.* **2018**, *48*, 87–95. [[CrossRef](#)]
37. Munko, M.J.; Rushton, L.M.; Ellis, L.M.; Zipfel, J.D.; Bevington, P.; Chalupczak, W.; Lopez Dubon, S. Sparse Sensing and Machine Learning for Rapid Calibration of Defect Detection with rf Atomic Magnetometers. *Sensors* **2025**, *25*, 6930. [[CrossRef](#)]
38. Candès, E.J.; Li, X.; Ma, Y.; Wright, J. Robust principal component analysis? *J. ACM* **2011**, *58*, 11. [[CrossRef](#)]
39. Lin, Z.; Chen, M.; Ma, Y. The Augmented Lagrange Multiplier Method for Exact Recovery of Corrupted Low-Rank Matrices. *arXiv* **2013**, arXiv:1009.5055. [[CrossRef](#)]
40. Scherl, I.; Strom, B.; Shang, J.K.; Williams, O.; Polagye, B.L.; Brunton, S.L. Robust principal component analysis for modal decomposition of corrupt fluid flows. *Phys. Rev. Fluids* **2020**, *5*, 054401. [[CrossRef](#)]
41. Sid-Lakhdar, W.; Cayrols, S.; Bielich, D.; Abdelfattah, A.; Luszczek, P.; Gates, M.; Tomov, S.; Johansen, H.; Williams-Young, D.; Davis, T.; et al. PAQR: Pivoting Avoiding QR factorization. In *Proceedings of the 2023 IEEE International Parallel and Distributed Processing Symposium (IPDPS), St. Petersburg, FL, USA, 15–19 May 2023*; IEEE: New York, NY, USA, 2023; pp. 322–332. [[CrossRef](#)]
42. Bijmens, L.; Verbeeck, R.; Göhlmann, H.; Talloen, W.; Ion, R.; Lewi, P.; Wouters, L. Spectral Map Analysis of Microarray Data. In *Comprehensive Chemometrics*; Elsevier: Amsterdam, The Netherlands, 2009; pp. 569–583. [[CrossRef](#)]
43. Pazur, K.; Bogusz, P.; Krasoń, W. Utilizing High-Speed 3D DIC for Displacement and Strain Measurement of Rotating Components. *Materials* **2025**, *18*, 3974. [[CrossRef](#)] [[PubMed](#)]
44. Huda, Z. *Mechanical Behavior of Materials*, 1st ed.; Mechanical Engineering Series; Springer: Cham, Switzerland, 2022. [[CrossRef](#)]
45. Yang, B. Stress analysis in two-dimensional problems. In *Stress, Strain, and Structural Dynamics*; Elsevier: Amsterdam, The Netherlands, 2023; pp. 155–179. [[CrossRef](#)]
46. Siebert, T.; Hack, E.; Lampeas, G.E.A. Uncertainty Quantification for DIC Displacement Measurements in Industrial Environments. *Exp. Tech.* **2021**, *45*, 685–694. [[CrossRef](#)]
47. Reu, P.L. *The story of DIC Uncertainty Quantification, Thus Far...*; Sandia National Laboratories: Albuquerque, NM, USA, 2017.
48. MatchID. [Software]. 2024. Available online: <https://www.matchid.eu> (accessed on 8 August 2025).
49. Aja-Fernández, S.; Ramos-Llordén, G.; Yushkevich, P.A. Image representation and 2D signal processing. In *Medical Image Analysis*; Elsevier: Amsterdam, The Netherlands, 2024; pp. 115–143. [[CrossRef](#)]
50. Gao, S.; Gruev, V. Bilinear and bicubic interpolation methods for division of focal plane polarimeters. *Opt. Express* **2011**, *19*, 26161. [[CrossRef](#)]
51. Anand, A.; Singh, P.; Srivastava, P.K.; Gupta, M. Chapter 19—GIS-based analysis for soil moisture estimation via kriging with external drift. In *Agricultural Water Management*; Srivastava, P.K., Gupta, M., Tsakiris, G., Quinn, N.W., Eds.; Academic Press: Cambridge, MA, USA, 2021; pp. 391–408. [[CrossRef](#)]
52. Oliver, M.; Webster, R. A tutorial guide to geostatistics: Computing and modelling variograms and kriging. *Catena* **2014**, *113*, 56–69. [[CrossRef](#)]

53. Hilal, A.; Bangroo, S.A.; Kirmani, N.A.; Wani, J.A.; Biswas, A.; Bhat, M.I.; Farooq, K.; Bashir, O.; Shah, T.I. Chapter 19—Geostatistical modeling—A tool for predictive soil mapping. In *Remote Sensing in Precision Agriculture*; Lamine, S., Srivastava, P.K., Kayad, A., Muñoz-Arriola, F., Pandey, P.C., Eds.; Earth Observation; Academic Press: Cambridge, MA, USA, 2024; pp. 389–418. [[CrossRef](#)]
54. Behrens, T.; Rossel, R.A.V.; Kerry, R.; MacMillan, R.; Schmidt, K.; Lee, J.; Scholten, T.; Zhu, A.X. The Relevant Range of Scales for Multi-Scale Contextual Spatial Modelling. *Sci. Rep.* **2019**, *9*, 14800. [[CrossRef](#)]
55. Damelin, S.B.; Hoang, N.S. On Surface Completion and Image Inpainting by Biharmonic Functions: Numerical Aspects. *Int. J. Math. Math. Sci.* **2018**, *2018*, 3950312. [[CrossRef](#)]
56. Zhou, H.F.; Dou, H.S.; Qin, L. A review of full-scale structural testing of wind turbine blades. *Renew. Sustain. Energy Rev.* **2014**, *33*, 177–187. [[CrossRef](#)]
57. Al-Khudairi, O.; Hadavinia, H.; Little, C.; Gillmore, G. Full-Scale Fatigue Testing of a Wind Turbine Blade in Flapwise Direction and Examining the Effect of Crack Propagation on the Blade Performance. *Materials* **2017**, *10*, 1152. [[CrossRef](#)]
58. Lee, H.G.; Lee, J. Measurement theory of test bending moments for resonance-type fatigue testing of a full-scale wind turbine blade. *Compos. Struct.* **2018**, *200*, 306–312. [[CrossRef](#)]
59. Blaber, J.; Adair, B.; Antoniou, A. Ncorr: Open-Source 2D Digital Image Correlation Matlab Software. *Exp. Mech.* **2015**, *55*, 1105–1122. [[CrossRef](#)]
60. Wang, Y.; Lava, P.; Reu, P.; Debruyne, D. Theoretical Analysis on the Measurement Errors of Local 2D DIC: Part II Assessment of Strain Errors of the Local Smoothing Method—Approaching an Answer to the Overlap Question. *Strain* **2016**, *52*, 129–147. [[CrossRef](#)]
61. Hild, F.; Roux, S. Digital Image Correlation: From Displacement Measurement to Identification of Elastic Properties—A Review. *Strain* **2006**, *42*, 69–80. [[CrossRef](#)]
62. Schreier, H.; Braasch, J.R.; Sutton, M.A. Systematic errors in digital image correlation caused by intensity interpolation. *Opt. Eng.* **2000**, *39*, 2915–2921. [[CrossRef](#)]
63. Blaysat, B.; Grédiac, M.; Sur, F. Effect of interpolation on noise propagation from images to digital image correlation displacement maps. *Int. J. Numer. Methods Eng.* **2016**, *108*, 213–232. [[CrossRef](#)]
64. Pan, B.; Qian, K.; Xie, H.; Asundi, A. Two-dimensional digital image correlation for in-plane displacement and strain measurement: A review. *Meas. Sci. Technol.* **2009**, *20*, 062001. [[CrossRef](#)]
65. Amiot, F.; Bornert, M.; Doumalin, P.; Dupré, J.-C.; Fazzini, M.; Orteu, J.-J.; Poilâne, C.; Robert, L.; Rotinat, R.; Toussaint, E.; et al. Assessment of Digital Image Correlation Measurement Accuracy in the Ultimate Error Regime: Main Results of a Collaborative Benchmark. *Strain* **2013**, *49*, 483–496. [[CrossRef](#)]
66. *DNVGL-ST-0164*; Tidal Turbines. DNV GL: Høvik, Norway, 2015.
67. Hoyez, H.; Schockaert, C.; Rambach, J.; Mirbach, B.; Stricker, D. Unsupervised Image-to-Image Translation: A Review. *Sensors* **2022**, *22*, 8540. [[CrossRef](#)] [[PubMed](#)]
68. Nyamathulla, S.; Veeranjanyulu, N. Analysis of Pix2Pix and CycleGAN for Image-to-Image Translation: A Comparative Study. In *Proceedings of the 2024 IEEE International Conference on Smart Power Control and Renewable Energy (ICSPCRE), Rourkela, India, 19–21 July 2024*; IEEE: New York, NY, USA, 2024; pp. 1–6. [[CrossRef](#)]
69. Isola, P.; Zhu, J.Y.; Zhou, T.; Efros, A.A. Image-to-Image Translation with Conditional Adversarial Networks. In *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 21–26 July 2017*; IEEE: New York, NY, USA, 2017; pp. 5967–5976. [[CrossRef](#)]
70. Sumera; Subashini, T.S.; Vaidehi, K. Realizing GAN Potential for Image Generation and Image-To-Image Translation Using Pix2Pix. In *Computer Vision and Image Processing, Proceedings of the CVIP 2024, Chennai, India, 19–21 December 2024*; Kakarla, J., Balasubramanian, R., Murala, S., Vipparthi, S.K., Gupta, D., Eds.; Communications in Computer and Information Science; Springer: Cham, Switzerland, 2026; Volume 2478, pp. 103–115. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.