

PYVALE: A Fast, Scalable, Open-Source 2D Digital Image Correlation (DIC) Engine Capable of Handling Gigapixel Images

Joel Hirst^{1*}, Lorna Sibson^{1,2}, Adel Tayeb¹, Ben Poole¹, Megan Sampson¹, Wiera Bielajewa¹, Michael Atkinson¹, Alex Marsh¹, Rory Spencer¹, Robert Hamill¹, Cory Hamelin¹, Allan Harte¹, Lloyd Fletcher¹

¹United Kingdom Atomic Energy Authority, Culham Campus, Abingdon, OX14 3DB, Oxfordshire, United Kingdom.

²School of Electrical and Electronic Engineering, University of Sheffield, Mappin Street, Sheffield, S1 4DT, South Yorkshire, United Kingdom.

*Corresponding author(s). E-mail(s): joel.hirst@ukaea.uk;

Abstract

Background: Digital Image Correlation (DIC) is a widely used full-field measurement technique, but both open-source and commercial packages often have limitations such as operating-system restrictions, lack of support for deployment on computing clusters, and poor scalability to gigapixel-scale images common in Scanning Electron Microscopy DIC (SEM-DIC). **Objective:** PYVALE is an open-source software package designed for sensor simulation, uncertainty quantification, placement optimization, and calibration/validation. A key component of this is the development of a dedicated 2D DIC module intended for standalone use and integration within broader workflows. **Methods:** PYVALE provides a user-friendly Python interface with performant compiled routines underneath. At its core is a multithreaded, reliability-guided DIC algorithm. Its open-source MIT license enables wide deployment, including on computing clusters and in automated pipelines. **Results:** Benchmarking with the publicly available 2D DIC challenge 2.0 dataset shows that PYVALE achieves metrological performance comparable to existing commercial and open-source DIC codes. It can correlate gigapixel-scale image pairs in under 5 minutes on high-specification desktop workstations, with memory peaking at approximately 50 GB. **Conclusions:** PYVALE's strong metrological foundation, coupled with its scalability for SEM-DIC, positions it as a platform for sustained, community-driven development. Its design

and licensing provide a foundation for future improvements in open-source DIC and integration into experimental design and validation workflows.

Keywords: Digital Image Correlation, DIC, Metrology, Experimental Mechanics, Full-field Measurement, Computational Mechanics, Open-Source Software

1 Introduction

Digital Image Correlation (DIC) is a widely used technique for measuring shape deformation and motion [1–3]. The idea is to obtain displacements by correlating gray-level values between a reference image and one or more images captured post deformation. In local subset DIC, a non-linear optimization routine is used to calculate the shape function (rigid, affine, quadratic etc.) parameters that describe the displacement and deformation of the subset between reference and deformed images. A single-camera setup (2D-DIC) enables measurement of in-plane displacements only, while stereo configurations extend this to capture out-of-plane motion. Since the first use of DIC in experimental mechanics by Peters and Ranson in 1982 [4], the field has seen vast improvements in the underlying computational and numerical algorithms. Some key algorithmic developments include the use of the Newton-Raphson (NR) method for minimizing cost functions [5], using fast-Fourier transforms (FFTs) for integer pixel displacement estimation [6], Reliability-Guided DIC (RG-DIC) [7], and understanding systematic errors due to image interpolation at the gray-level [8].

These methodological advancements have, in turn, led to the creation of several well-established software packages. Commercial codes include MatchID [9], DaVis [10], VIC-2D/VIC-3D [11] to name a few, while open-source local subset DIC codes, such as Ncorr [12], OpenCorr [13], and DICe [14], are widely used in academic and research settings. Ncorr comes as a MATLAB extension with an easy installation process and intuitive interactive Graphical User Interface (GUI) for correlation configuration and Region of Interest (ROI) selection. However, there is very little scripting support for the extension, making systematic studies with correlation parameters a challenge. OpenCorr and DICe are both C++ codes that provide prebuilt executables for Windows. For Linux operating systems, the package and dependencies must be built from source which makes installation a time consuming process. Another notable mention goes to μ DIC [15], which is a well documented, open-source global DIC code. μ DIC is written purely in Python which, while making usage straightforward, means it struggles computationally with higher resolution images.

Commercial tools are built with user friendliness in mind, offering intuitive GUIs and reliable performance. However, they often come with significant financial cost and limited flexibility, with users typically unable to access or modify the underlying algorithms beyond what the vendor is willing to implement or permit. Another key limitation of commercial codes is licensing that restricts the ability of users to perform large distributed array jobs on computing clusters for experimental design and uncertainty quantification. Open-source alternatives offer greater adaptability, allowing users to modify and extend the software to suit specific needs. These open-source codes also have the benefit that they can be run without licensing restrictions on computing clusters. However, this flexibility often comes at the expense of usability or performance. DIC codes written in commonly used data analysis languages, such as Python, can be user-friendly but struggle with memory usage and computational performance. In contrast, lower-level compiled languages (e.g., C/C++) offer improved performance and memory management but require more complex setup and

dependency management, beyond the typical capabilities of engineers not focused on computational work. Furthermore, open-source tools that rely on proprietary environments such as MATLAB still require a paid license, which hinders accessibility. Therefore, there is a gap in available open-source DIC software that combines an intuitive and easy-to-use interface with the computational performance of compiled code through Python C extensions.

To study material deformation at the microscopic scale, techniques that exceed the resolution limits of optical microscopy (~ 200 nm) are essential. However, a major limitation of current open-source DIC tools is their inability to efficiently handle extremely large images, such as gigapixel-scale datasets required to capture representative areas with displacement resolution to reveal micro-mechanics modes of deformation from Scanning Electron Microscope (SEM) DIC. Electron-based imaging enables analysis at much smaller scales due to the shorter wavelength of electrons. SEM-DIC, particularly using field emission SEMs (FESEM), has gained traction in recent years [16–22], though recent work explores more cost effective systems that use tungsten or LaB_6 emitters [23]. A key challenge with SEM-DIC is the speckle pattern generation. As per the DIC good practices guide, the ideal size of each speckle should be between 3-5 pixels [3]. At SEM resolutions this translates to a physical speckle size in the nanometer range. An in-depth review of methods for speckle pattern generation at these scales can be found in Ref. [20]. The fine spatial resolution combined with a large ROI can still produce comparatively large subset displacements in pixel units. It is often the case that with an optimal speckle pattern size of 3-5 pixels, subset displacements can be well in excess of 1000 pixels. To capture these large deformations while maintaining sufficiently high spatial resolution, SEM captures are stitched together. With images approaching gigapixel sizes, memory management and processing speed quickly become one of the most important factors when selecting a DIC code.

To address this, we have designed the 2D DIC module in PYVALE to meet the following criteria: (i) open-source with permissive licensing allowing scalable usage for large array jobs on the super-computing clusters, (ii) a simple and accessible Python based user interface, (iii) memory efficient, parallelized low-level code for good computational performance (written in C/C++), (iv) straightforward installation across all major operating systems and computing architectures, and (v) the ability to process gigapixel size images for SEM-DIC applications. This gives PYVALE users the simplicity of a Python interface combined with the performance of optimized compiled code, all under a free and open-source license that supports deployment on computing clusters.

Here we will present an in-depth overview of PYVALE’s user interface, the algorithms used, as well as benchmarking current metrological and computational performance. The paper is organized as follows: Section 2 presents the design and implementation of PYVALE, including the user interface and key algorithms. Section 3 provides numerical validation against the 2D DIC challenge. Section 4 evaluates performance

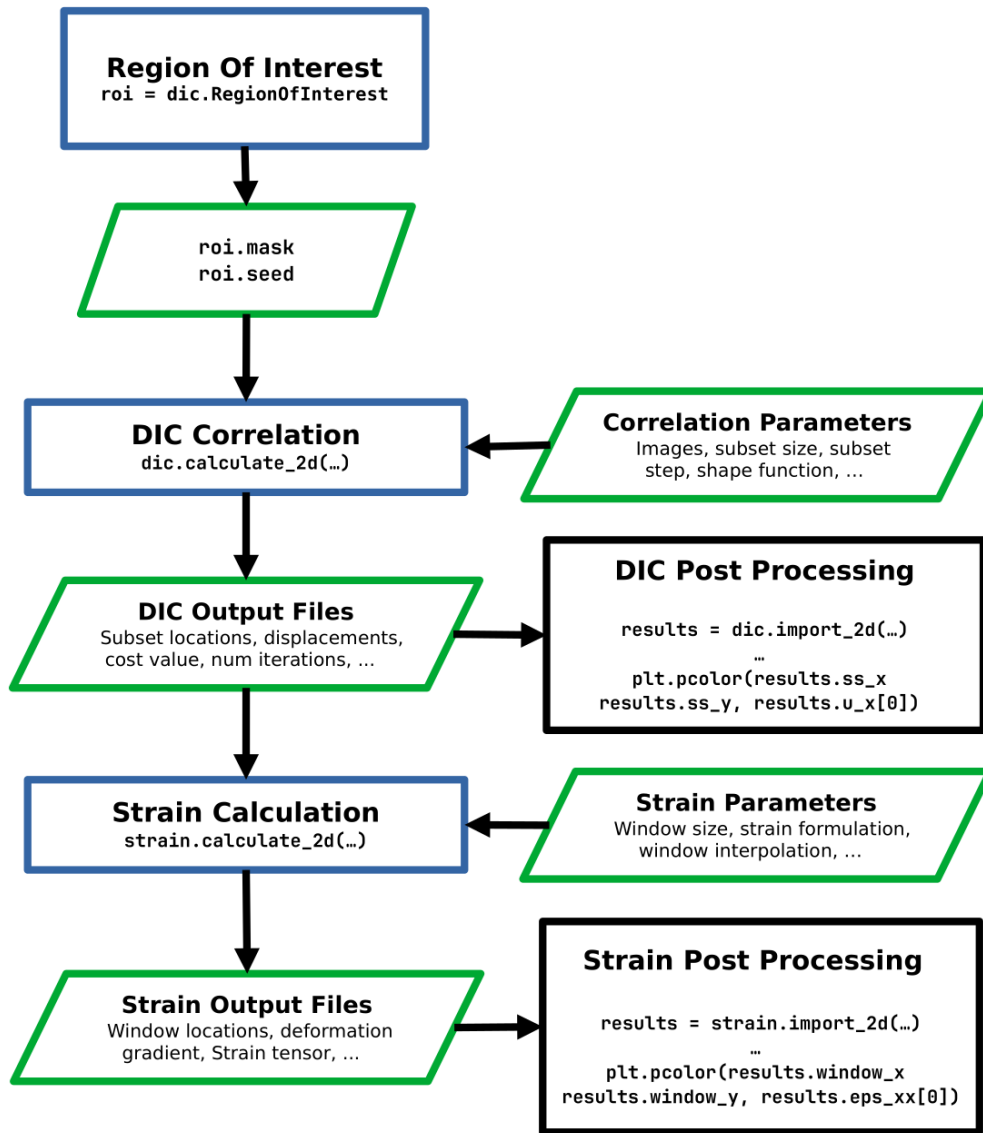


Fig. 1 High-level flowchart illustrating the PYVALE workflow. Each box shows the required core functions, parameters, and arguments. The two post-processing boxes give examples of how data can be plotted with `matplotlib`, although the data can also be imported into other analysis packages (e.g., Excel or Origin) at user discretion.

on gigapixel SEM-DIC images, with a focus on runtime and memory scaling. Section 5 presents a runtime comparison to other open-source codes using both SEM-DIC and synthetic speckle pattern images. Finally, section 6 provides an outlook on future DIC development and user engagement.

2 Design & Implementation

2.1 User Interface & Workflow

At a high level, PYVALE’s DIC engine is a Python module that uses Python C extensions to make calls down to lower level C/C++ code for computationally intensive calculations. When installed via Python’s `pip` package manager on major operating systems (Windows 10/11, macOS, Linux), the C/C++ components are precompiled which significantly simplifies the setup process. Typical users are not required to configure or install multiple dependencies manually; instead, they receive a ready-to-use package for their system. For users needing greater control or wishing to make modifications, PYVALE also supports manual compilation from source.

A flowchart summarizing key components of the PYVALE DIC workflow from a user’s perspective is shown in Fig. 1. To run PYVALE’s DIC engine with the minimum required amount of input, the user must specify the ROI over which they want to perform the correlation and some fundamental correlation parameters such as subset-size and subset-step. The ROI selection can be completed in one of three ways: (i) interactively via a built-in GUI tool, (ii) by defining a configuration in a “.yaml” file, or (iii) by creating a boolean array either manually (for example using standard NumPy array operations) or using the helper functions included in PYVALE. The ability to programmatically manipulate the ROI without having to open images is particularly useful for larger resolutions (i.e. gigapixel-scale), which can cause lag/crashes when opened within a GUI.

Once an ROI has been chosen, the user can proceed with a numerical correlation to obtain subset displacement values. A simple DIC calculation example can be found in Listing 1. Here, the user sets the reference and deformed images prior to calling the interactive ROI selection tool. The DIC analysis is then performed using a subset-size of 31 pixels and step of 5 pixels, with the ROI mask and starting seed location selected during the interactive ROI session. For each deformed image, there will be a corresponding result file containing correlation data. There are a wide range of additional arguments that can be passed to the DIC engine, allowing the user to tailor the setup to suit their needs. These include the ability to manipulate the convergence conditions, output file formatting, number of threads. Further details can be found in the package documentation [24]. By default, results are saved in human-readable comma-separated value (CSV) format in the current working directory. Each file will be prefixed with “dic_results_” followed by the original name of the deformed image unless otherwise specified. The delimiter, output folder and filename prefix can be altered if required. Data can also be saved in a non-human-readable binary format, which is particularly useful for large images where writing to human-readable text format becomes a time-consuming part of the workflow. In both cases, PYVALE provides the ability to import results into a structured dataclass composed of NumPy arrays for custom analysis and visualization.

```

1 import pyvale.dic as dic
2
3 ref_image="./ref_0000.tiff"
4 def_image="./def_*.tiff"
5
6 roi = dic.RegionOfInterest(ref_image)
7 roi.interactive_selection()
8
9 dic.calculate_2d(reference=ref_image,
10                deformed=def_image,
11                roi_mask=roi.mask,
12                seed=roi.seed,
13                subset_size=31,
14                subset_step=15,
15                output_basepath="./")

```

Listing 1 A simple example using the interactive ROI selection and 2D DIC in PYVALE. Accurate as of PYVALE release 2026.1.1.

```

1 import matplotlib.pyplot as plt
2 import pyvale.dic as dic
3
4 dic_data = dic.import_2d(data="./dic_results_*.tiff",
5                         binary=False,
6                         delimiter=",")
7
8 # plot of vertical displacement for first deformation image.
9 plt.pcolor(dic_data.ss_x,
10           dic_data.ss_y,
11           dic_data.u_y[0]) # [image, y, x]
12 plt.show()

```

Listing 2 Example of importing PYVALE DIC displacement data into Python and plotting the vertical displacement field for deformed image 0. Accurate as of PYVALE release 2026.1.1.

While PYVALE does not currently include built-in plotting functionality for displacement and strain data, many Python packages (e.g., Matplotlib [25], Plotly [26]) can be used as part of the same script to visualize results. A simple example demonstrating how to import DIC data and generate a 2D plot of the vertical displacement in the first deformed image is shown in Listing 2.

PYVALE DIC result files contain displacement fields and can be used to perform a deformation gradient and strain field calculation. A user can pass the type of strain formulation to use, the strain window size, whether to use bilinear or biquadratic interpolation on the window, and options for output data file formatting. Strain data can be imported in a similar way to DIC data. An example showing a simple strain calculation and importing the results can be found in Listing 3.

2.2 Underlying DIC Algorithms

The capability of PYVALE to handle gigapixel-scale images and large displacement fields is attributed to a two-stage processing approach. Initially, a multi-window FFT

```

1 import matplotlib.pyplot as plt
2 import pyvale.dic as dic
3 import pyvale.strain as strain
4
5 strain.calculate_2d(data="./dic_results_*",
6                   input_binary=False,
7                   input_delimiter=",",
8                   window_size=5,
9                   window_element=9,
10                  output_basepath="./")
11
12 strain_data = strain.import_2d(data="strain_*",
13                              delimiter=",")
14
15 # plot xx component of 2D strain tensor for
16 # first deformed image.
17 plt.pcolor(strain_data.window_x,
18           strain_data.window_y,
19           strain_data.eps_xx[0]) # [image, y, x]
20
21 plt.show()

```

Listing 3 An example of performing a strain calculation and plotting the resulting horizontal strain field. Accurate as of PYVALE release 2026.1.1.

based technique, which is a key method in Particle Image Velocimetry (PIV) [27–29], is applied to estimate rigid displacements for all subsets within the user-specified ROI. Similar methods have been used in previous versions of DaVis [10]. The DaVis implementation of the multi-window FFT method has previously been benchmarked against other commercial and open-source codes specifically for SEM-based micro-mechanics applications in [30]. The multi-window procedure begins by splitting images into square subsets initially with dimensions larger than the user estimated maximum displacement. Within each window, a normalized FFT-based cross-correlation (FFTCC) is computed to estimate the rigid-body displacement. The open-source PocketFFT library [31] is used for all FFT calculations in PYVALE. The FFT-based normalized cross-correlation is given by:

$$R_{fg}(x, y) = \mathcal{F}^{-1} \left\{ \frac{\mathcal{F}\{f(x, y)\} \overline{\mathcal{F}\{g(x, y)\}}}{|\mathcal{F}\{f(x, y)\}| |\mathcal{F}\{g(x, y)\}|} \right\}, \quad (1)$$

where \mathcal{F} and \mathcal{F}^{-1} denote the forward and inverse 2D Fourier transform respectively, $\overline{\mathcal{F}\{g(x, y)\}}$ is the complex conjugate of the Fourier transform of $g(x, y)$, where $f(x, y)$ is the reference window, and $g(x, y)$ is the corresponding window in the deformed image. The location of the peak in the cross-correlation map $R_{fg}(x, y)$ yields the rigid displacement vector (u, v) for that particular window. After estimating the rigid shift, the window size is halved. The displacements, (u_x, u_y) , from the four nearest points in the previous, coarser window are then used as the initial guess for the rigid displacement in the newly halved window. This process of halving the window size continues iteratively until the user-defined subset size is reached. An illustration of this procedure can be found in the left-hand side of Fig 2.

PYVALE determines window sizes automatically (excluding the user specified subset size) by selecting powers of two (2^n) descending from the next highest power of two above the user estimated maximum displacement. This exploits the inherently faster FFT computations for powers of two compared to arbitrary array lengths [32]. As an example of this, take the case where a user estimates that the maximum displacement is approximately 800 pixels and specifies a subset-size of 17 pixels. The PYVALE algorithm will use window sizes of 1024, 512, 256, 128, 64, 32 and finally 17 to estimate the rigid displacement. The rigid values of the final window (in this case 17) will be used as seed values for the RG-DIC algorithm [7]. For micro-mechanics applications it may be desirable to only run the multi-window FFT correlation without a subsequent RG-DIC step as the contribution of the higher order shape function parameters can be small. In this case the user can set `method='MULTIWINDOW'` when calling `dic.calculate_2d`. To obtain subpixel displacement values using only the multi-window approach, a 2D gaussian is fitted in the vicinity of the correlation peak in Fourier space. Generally, it is recommended that the user use `method='MULTIWINDOW_RG'` (default), even for rigid displacements. A bicubic b-spline is used to obtain intensity values at subpixel locations.

For each subset, displacement is obtained by minimizing a cost function. The Sum of Squared Differences (SSD) compares gray-level intensities between the reference and deformed images:

$$\text{SSD} = \sum_i (f(x_i, y_i) - g(x_i, y_i))^2 \quad (2)$$

where $f(x_i, y_i)$ and $g(x_i, y_i)$ represent the gray-level intensity values at location (x_i, y_i) in the reference and deformed images, respectively. To reduce sensitivity to intensity scaling, the Normalized SSD (NSSD) is defined as:

$$\text{NSSD} = \sum_i \left(\frac{f(x_i, y_i)}{\sqrt{\sum_j f(x_j, y_j)^2}} - \frac{g(x_i, y_i)}{\sqrt{\sum_j g(x_j, y_j)^2}} \right)^2 \quad (3)$$

The Zero-Normalized SSD (ZNSSD), used by default in PYVALE, further accounts for intensity offsets:

$$\text{ZNSSD} = \sum_i \left(\frac{\bar{f}(x_i, y_i)}{\sqrt{\sum_j \bar{f}(x_j, y_j)^2}} - \frac{\bar{g}(x_i, y_i)}{\sqrt{\sum_j \bar{g}(x_j, y_j)^2}} \right)^2 \quad (4)$$

where $\bar{f}(x_i, y_i) = f(x_i, y_i) - f_m$, and f_m is the mean gray-level value of the subset. ZNSSD is invariant to both additive and multiplicative intensity changes. There are three options for shape functions that transforms pixel coordinates in the subset from

the reference image to the deformed image:

$$\xi(x_i, y_i, \mathbf{P}) = \underbrace{\begin{bmatrix} p_0 \\ p_1 \end{bmatrix}}_{\text{rigid}} + \underbrace{\begin{bmatrix} 1 + p_2 & p_3 \\ p_4 & 1 + p_5 \end{bmatrix} \begin{bmatrix} x_i \\ y_i \end{bmatrix}}_{\text{affine}} \dots \quad (5)$$

$$\dots + \underbrace{\begin{bmatrix} p_6 & p_7 & p_8 \\ p_9 & p_{10} & p_{11} \end{bmatrix} \begin{bmatrix} x_i^2 \\ x_i y_i \\ y_i^2 \end{bmatrix}}_{\text{quadratic}}$$

Each higher-order shape function includes all terms from the lower-order functions: affine includes rigid terms, and quadratic includes both affine and rigid terms.

For the non-linear optimization for each subset a Levenberg-Marquardt [33, 34] algorithm is used. The convergence threshold and precision, as well as the maximum number of iterations for a given subset optimization can be altered by the user if needed. Default values are used when they are omitted by the user. There is a large amount of existing literature on optimizations of cost and shape functions in DIC and therefore the explanations are omitted here. Well explained mathematical descriptions can be found across various works [1, 2, 12, 35].

2.3 Strain Calculation and Formulations

From the displacement fields, the in-plane strain components can be calculated using the spatial derivatives of the displacement field to capture local changes in geometry. The 2D deformation gradient matrix, \mathbf{F} , is given by the equation:

$$\mathbf{F} = \begin{bmatrix} 1 + \frac{\partial u_x}{\partial x} & \frac{\partial u_x}{\partial y} \\ \frac{\partial u_y}{\partial x} & 1 + \frac{\partial u_y}{\partial y} \end{bmatrix} = \mathbf{I} + \nabla \mathbf{u}, \quad \mathbf{u} = \begin{bmatrix} u_x \\ u_y \end{bmatrix}. \quad (6)$$

To calculate the partial derivatives above, we calculate the gradient over a square window containing $N \times N$ displacement data points from the DIC calculation. Because we are using displacement values from DIC, the strain window is dependent on the subset-step, s , and subset-size, w . These quantities, along with the size of the strain window, form what is typically referred to as the Virtual Strain Gauge, VSG :

$$VSG = (N - 1)s + w \quad (7)$$

Due to the noise in DIC measurements, smoothing is typically applied over the strain window. PYVALE supports bilinear and biquadratic smoothing over the strain window elements. The polynomial approximation is given by:

$$\mathbf{u}(x, y) = \mathbf{P}(x, y) \mathbf{c}, \quad (8)$$

where $\mathbf{u} = [u_x \ u_y]^T$, $\mathbf{P}(x, y)$ is a row vector of basis terms, and \mathbf{c} is a column vector of coefficients. The polynomial basis is

$$\mathbf{P}(x, y) = \begin{cases} [1, x, y] & \text{bilinear,} \\ [1, x, y, x^2, y^2, x^2y, xy^2, x^2y^2] & \text{biquadratic} \end{cases} \quad (9)$$

The coefficients are obtained by solving a linear least-squares problem, since the displacement field is linearly parameterized in the unknown coefficients. Once the polynomial coefficients have been obtained the deformation gradient tensor and strain can be calculated. PYVALE supports the following strain tensor formulations:

- *Green-Lagrange*: $\boldsymbol{\varepsilon} = \frac{1}{2} (\mathbf{F}^T \mathbf{F} - \mathbf{I})$
- *Hencky (logarithmic)*: $\boldsymbol{\varepsilon} = \ln(\sqrt{\mathbf{F}^T \mathbf{F}})$
- *Euler-Almansi*: $\boldsymbol{\varepsilon} = \frac{1}{2} (\mathbf{I} - (\mathbf{F} \mathbf{F}^T)^{-1})$
- *Biot (right / Lagrangian)*: $\boldsymbol{\varepsilon} = \sqrt{\mathbf{F}^T \mathbf{F}} - \mathbf{I}$
- *Biot (left / Eulerian)*: $\boldsymbol{\varepsilon} = \sqrt{\mathbf{F} \mathbf{F}^T} - \mathbf{I}$

where \mathbf{I} the 2×2 identity matrix.

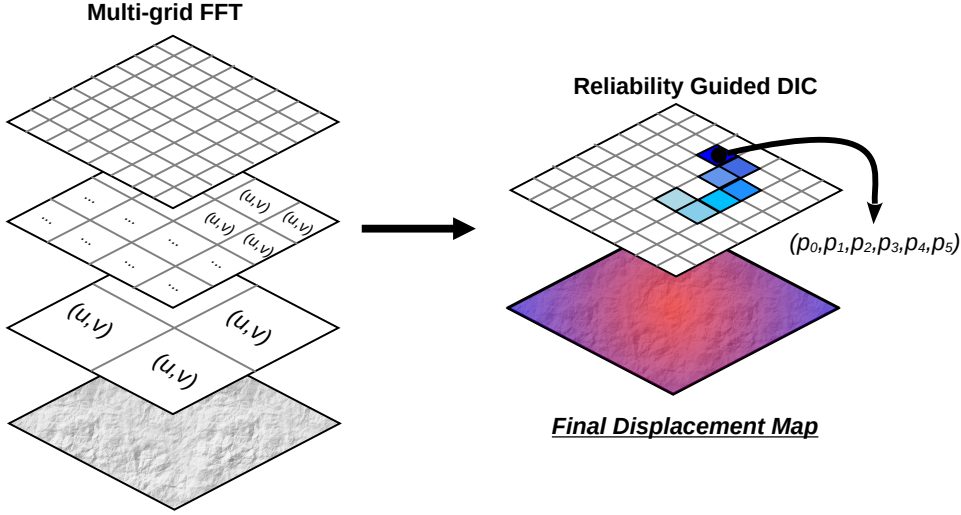


Fig. 2 Schematic of the main algorithmic approach taken in Pyvale. The first stage is to perform a FFT cross correlation to obtain rigid displacement values for decreasing window size. The final window size is then used to seed the RG-DIC to obtain higher order shape function parameters if required.

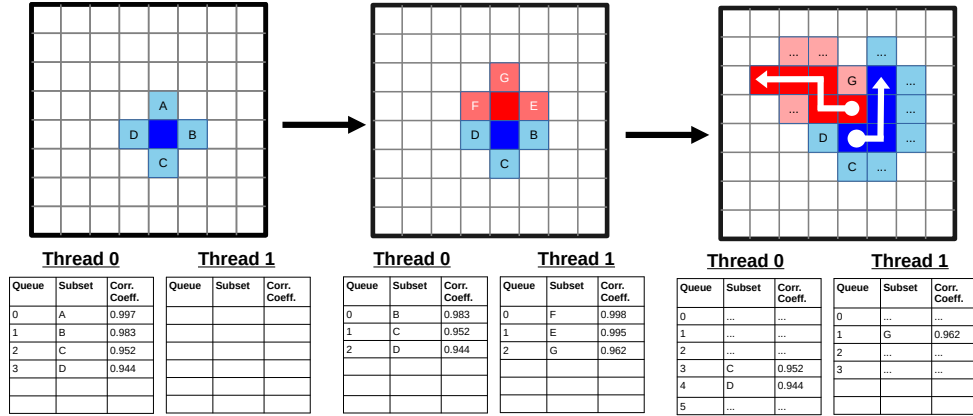


Fig. 3 Schematic of the multithreaded RG-DIC algorithm adopted in PYVALE. Here it is shown for two threads. (Left) A single thread performs calculation for the seed point and its 4 neighbours. (Middle) The second thread then “steals” work from the queue of thread 0 to start its own RG-DIC procedure and computes correlation values for the “stolen” point. These are then added to its own queue. (Right) A view of the queue sometime later where each thread has gone off in independent directions, note the position of subsets in the queue have changed depending on the ranking of the correlation coefficient.

2.4 Parallelization

PYVALE uses the Open Multi-Processing (OpenMP) library [36] to parallelize all major computations, including subset generation, bicubic interpolation setup, multi-window FFT-based rigid displacement estimation, and the RG-DIC algorithm. The multi-window FFT-displacement estimation is parallelized by evenly distributing windows among threads, with each thread storing rigid displacement values in a shared array to be used as initial estimates for subsequent iterations.

PYVALE uses a dynamic queuing parallelization strategy, which is contrasted against the thread-specific RG-DIC seed assignment approach used in the MATLAB extension NCorr [12]. An example schematic of the queuing process for the PYVALE multi-threaded RG-DIC is illustrated in Fig. 3. In our method, we start with a single user specified seed location from which a single thread then calculates the correlation coefficient for this point and its four nearest neighbors. If the seed location and its neighbors successfully converge, then the neighboring points are added to a queue. This is a standard initial step in the RG-DIC algorithm. Where the algorithm of PYVALE differs is that other threads then look for *work* by checking the queues of the other threads and *stealing* the points from other threads if there are idle subsets waiting in the queue. Other threads always steal the point at the top of the other queues (highest correlation coefficient). This is then repeated until all threads have been allocated work in their own queue. Each thread then populates its own queue and continues along a thread specific RG-DIC path, checking and updating a global mask to prevent repeat calculations for individual subsets. Threads can run out of




Fig. 4 Example showing the displacement calculated using PYVALE’s 2D DIC engine for the *Star5* reference and deformed images from the DIC Challenge 2.0 dataset. The deformed image has a maximum ground truth displacement of ± 0.5 pixels. The convergence of the DIC displacement data along the horizontal dotted white line at the vertical midpoint to the ground truth value of 0.5 px is used to determine metrological performance. Images can be downloaded from <https://idics.org/challenge/>.

work - perhaps by hitting an ROI boundary or arriving at a point where all neighboring points have been calculated. In such cases, the thread will restart the process of searching for work in at the top of the queues of other active threads. If work is available, the thread will once again steal work and begin populating its own queue. This process repeats until correlation values have been computed across the entire ROI.

3 Case Study: DIC Challenge 2.0

The Society of Experimental Mechanics (SEM) together with the International DIC (iDICs) society put together the DIC challenge to compare performance and accuracy of DIC codes to a set of well-defined standards (<https://idics.org/challenge/>). At the time of writing, there have been two iterations of the 2D DIC challenge [37, 38], as well as stereo [39] and Digital Volume Correlation (DVC) [40] comparisons. In this section, we shall compare the metrological performance of PYVALE to other local subset-based DIC codes that participated in the second iteration of the 2D DIC challenge. A short description of the benchmark images from the challenge are provided below. Those interested in a full in-depth description of the participants, methods and image generation procedures can find further details through Ref [38].

The benchmark dataset created as part of the 2D DIC challenge 2.0 contains a series of images with a sinusoidal vertical displacement pattern with decreasing spatial frequency from left to right. In the DIC challenge paper [38], these are referred to as *star* images. The same naming convention will be applied here. We present a comparison of PYVALE to other codes for *Star5*. For this pattern, three reference images are provided: a noise free reference image, a reference image + noise, and a noisy deformed image. Each image is 4000×500 pixels with a maximum ground truth displacement of ± 0.5 pixels at peak amplitude. An example of the star pattern as calculated from PYVALE’s 2D DIC engine overlaid over the noisy reference image is shown in Fig. 4.

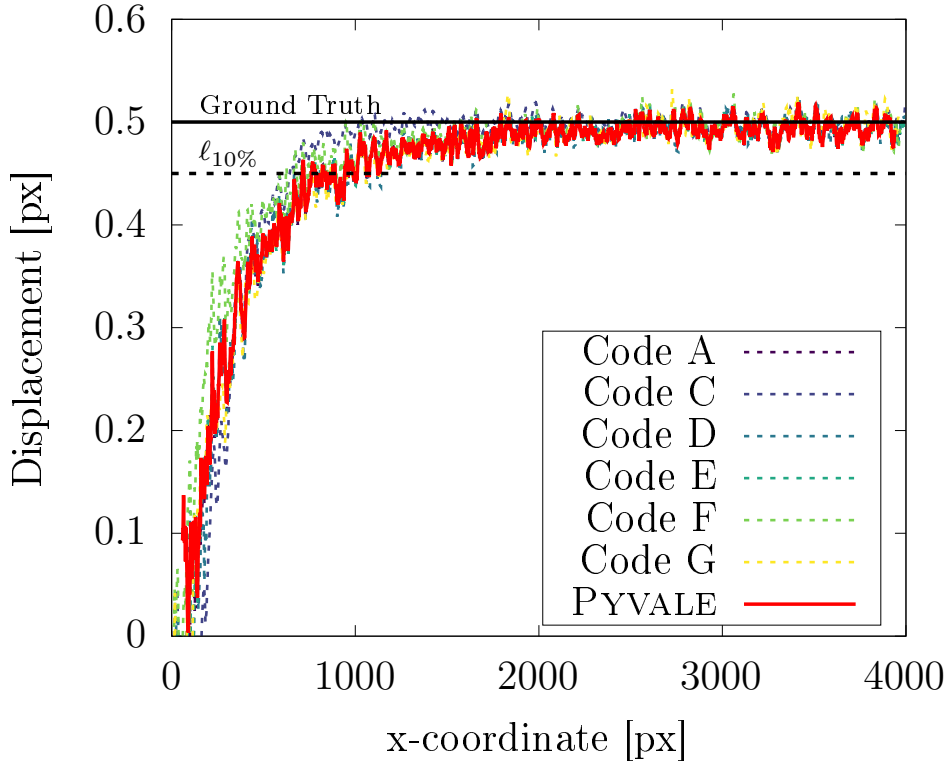


Fig. 5 A example comparison of the convergence of PYVALE (Solid red line) and along the midpoint horizontal (shown by the dotted white line in Fig. 4) anonymized DIC codes (dotted lines) to the ground truth displacement value of 0.5 px using the DIC challenge 2.0 images.

The assessment of metrological performance in the DIC challenge stems from the ability of a DIC code to converge to the ground truth displacement of 0.5 pixels along the horizontal direction for a fixed vertical displacement at $y = 250$ pixels (shown by the dotted white line in Fig. 4). An example comparison of the convergence to the ground truth segment between PYVALE and the anonymized codes can be found in Fig. 5. This example calculation was performed for a subset-size and step of 19 and 1 pixels respectively using an affine shape function with a ZNSSD cost function. PYVALE is shown by the solid red line, while all other anonymized codes are shown with dotted lines. Note that not all codes provided results for this specific subset-size, hence the omission of *Local.Affine.B* and *Local.Affine.H*. While this example is only for a single subset-size, PYVALE’s DIC engine can be seen to converge at an almost identical rate to other codes.

To fully assess the metrological performance of PYVALE we compare the spatial resolution, ℓ , and Metrological Efficiency Indicator, MEI , to the anonymised results from different local subset DIC codes. In the DIC challenge, an attenuation of 10%




Fig. 6 Spatial resolution of PYVALE (red circles) compared to other DIC codes (crosses) looked at as part of the DIC Challenge 2.0.

was selected as a reasonable cutoff to define the spatial resolution, $\ell_{10\%}$. This cutoff is shown by the dashed horizontal line in Fig. 5. A measure of $\ell_{10\%}$ was obtained by fitting a 12th order polynomial to midpoint convergences to the ground truth displacement for varying subset-size. From the fit, the period of oscillation (in pixels) at the x -coordinate where the fit crosses the 10% bias threshold defines $\ell_{10\%}$ for a given subset-size. A subset-step of 1 was used for all subset-sizes. The period of oscillation at the 10% cutoff for PYVALE alongside the anonymized codes from the DIC challenge can be found in Fig. 6. While PYVALE agrees well with most codes, there are two that stand out with a higher spatial resolution compared to the majority. As a result of the anonymization, we are unable to comment on the differences in implementations.

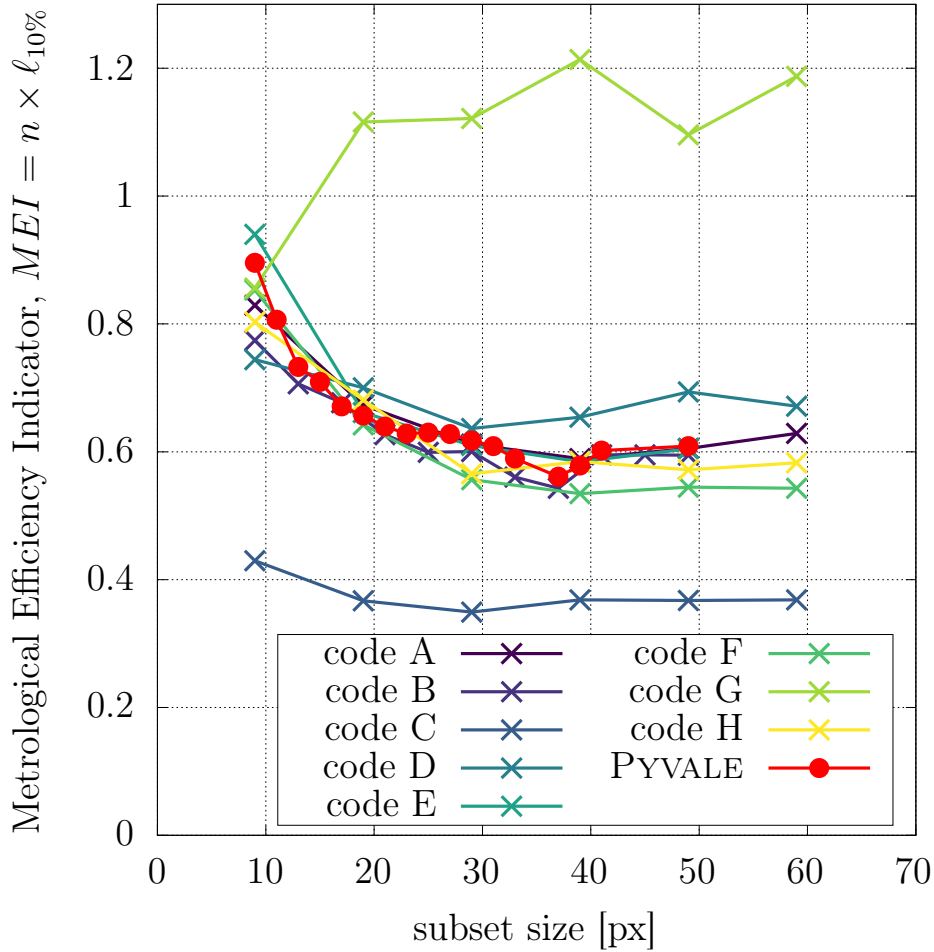


Fig. 7 Spatial resolution of PYVALE (red circles) compared to other DIC codes (crosses) looked at as part of the DIC Challenge 2.0.

Using results for the spatial resolution, a comparison to the MEI can be made. The MEI is defined by the measurement resolution multiplied by the spatial resolution [38, 41]:

$$MEI = n \times \ell_{10\%} \quad (10)$$

where n is the displacement noise ($n = 1\sigma$). This was obtained by performing a DIC calculation on the reference image against the reference image with added noise. Generally, the lower the MEI value, the better the metrological performance of a particular code. The MEI should theoretically remain constant as a function of subset-size, which has been analytically proven for Localized Spectrum Analysis (LSA) and empirically observed for both local and global DIC approaches[41, 42]. Fig. 7 shows a

Code	MEI	% Difference to Ref
A	0.602	+6.38%
B	0.566	Ref
C	0.361	-36.20%
D	0.654	+15.54%
E	0.600	+5.95%
F	0.541	-4.47%
G	1.023	+80.68%
H	0.573	+1.30%
PYVALE	0.576	+1.84%

Table 1 Average of the three lowest MEI values and percentage difference relative to code B (reference). Code B was used as a reference for the equivalent comparison in the DIC Challenge [38].

comparison between the MEI as a function of subset-size for PYVALE and the other DIC challenge codes. The final MEI value is calculated as the average of the three lowest values, which is inline with calculations in the DIC Challenge. A table of MEI values along with percentage differences to relative Code B can be found in Table 1.

This section has shown PYVALE has comparable metrological performance to other local-subset based DIC codes. Next we’ll discuss the capability to deal with extremely large SEM-DIC images (demonstrated in the next case study) and its computational efficiency compared to other open-source codes (demonstrated in the final case study).

4 Case Study: Application of Pyvale to Experimental Scanning Electron Microscope DIC (SEM-DIC) images

Here we will demonstrate PYVALE’s ability to handle gigapixel-scale SEM-DIC images using a reference and deformed experimental image pair of $32,000 \times 32,000$ pixels. The experiment was performed on oxygen-free high conductivity copper and captured using a Tescan Clara field emission scanning electron microscope, operated at 5 kV accelerating voltage, with a 300 pA probe current and a 3.2 μ s dwell time. To enable high resolution strain mapping over a large region of interest, a 19×19 array of $20 \mu\text{m} \times 20 \mu\text{m}$ tiles with a 2048×2048 image resolution were captured. A tile overlap of 20% was used to ensure accurate image stitching following capture. Image arrays were captured automatically using the NewTec Scientific Softstrain software, controlling the SEM through the API interface. Images were stitched using the ITK-Montage Python package [43]. The tensile specimen was deformed to a global strain of approximately 5% using a NewTec Scientific MT1000 microtensile tester.

We used the PYVALE 2D DIC engine to analyse the SEM-DIC images with a subset-size and step of 31 and 15 pixels respectively, a ZNSSD correlation function, and an affine shape function with bicubic b-spline interpolation for sub-pixel accuracy.

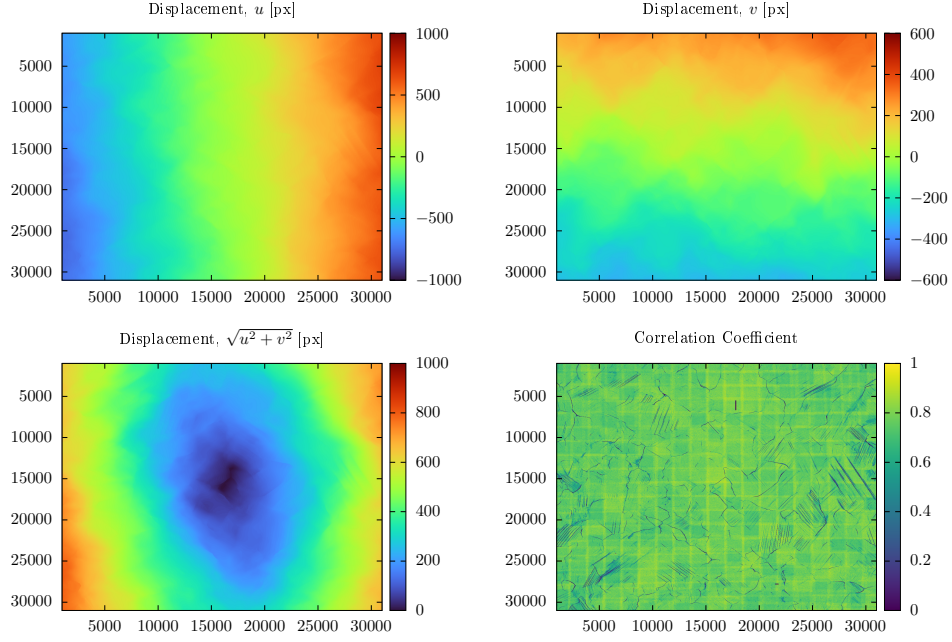


Fig. 8 Displacement values and correlation coefficient output from PYVALE for an example 32000×32000 image using the FFT + RG algorithm.

For each subset the maximum number of optimization iterations was set to 40 and convergence precision and convergence threshold set to 0.70 and 0.01 respectively. A Median Absolute Deviation (MAD) [44] outlier removal was applied to the estimated rigid displacements from the multi-window FFT approach. This was done at each window size iteration. In many cases, MAD outlier removal will prevent incorrect peak locations in the FFT correlation data propagating down through the window sizes. The ROI was created programmatically, with a 1000 pixel boundary region excluded along all outer edges of the reference image. No subset correlations were performed in this border region. Output from DIC calculations were saved to a human-readable file in a standard comma-separated value (CSV) format. The maximum displacement was set to 2000 pixels, which is well above the expected maximum of ~ 1000 pixels. The horizontal, vertical and magnitude displacement maps, along with a correlation coefficient map can be found in Fig. 8. The grid-like pattern that can be seen in the correlation coefficient map is a result of the stitching procedure used to combine the SEM data. Displacements have been saved for each and every subset in the correlation, hence the continuity in the displacement data. In many cases, a user may not want return displacement values for subsets that have not met the convergence criteria. In such a scenario a flag can be toggled to set all unconverged displacements to NaN values prior to saving to the results file. Further details can be found in the PYVALE documentation [24].

4.1 Runtime and Performance for SEM-DIC Images

To illustrate how PYVALE scales in both memory usage and runtime for images of this size, we examine both runtime scaling as a function of thread count on a single workstation, and memory scaling with variable image sizes for a fixed number of threads. The workstation used for these tests had an AMD Ryzen Threadripper 7980X (64 cores), 256 GB of DDR5 RAM, 1TB NVMe SSD, running Ubuntu 24.04.2 LTS (x86_64) as the operating system. Fig. 9 shows both the absolute runtime and the speedup relative to single-threaded performance for the SEM image pair. For the 32000×32000 pixel images analyzed here, processing with a single thread results in a correlation time of approximately 40 minutes. Using all 64 threads reduces this time to 48 seconds, nearly a 50-fold speedup. This is illustrated in the bottom panel of Fig. 9. The dotted line in the bottom pane shows an ideal scaling relationship with a one-to-one relationship between the relative speedup and increasing thread count.

In Fig. 9 PYVALE can be seen to drop below an ideal scaling relationship due to the growing amount of communication required between threads for the RG-DIC multi-threaded queuing. An alternative approach would be to have totally independent threads that work within an individual ROI (similar to Ncorr) where users must specify a seed location for each thread. This adds a significant amount of setup work for the user. Having to select all 128 seed locations and ensure convergence before running through the RG-DIC would be an arduous task. While increasing setup time, it is likely to lead to an increase in scaling performance if we assume converged seed locations as well as image defects, poor lighting, artifacts, etc. are distributed evenly between the individual thread ROIs. However, it is possible to envisage scenarios where the ROI associated with an individual thread covers a mostly artifact ridden portion of the overall ROI causing a specific thread to be stuck. Such cases would lead to an inflated runtime, with threads finishing their ROI at different times. Threads in “good” regions of the ROI will sit idle while threads covering “poor” regions will require a greater number of optimization iterations per subset.

Memory consumption and cache misses are important considerations for software performance, especially for the images used in this section. There are six main buffers in memory being responsible for most of the memory consumption in PYVALE’s 2D DIC implementation. Namely, these are the reference image, the stack of deformed images, directional derivatives used for bicubic interpolation, and a boolean mask for the ROI. Fig. 10 shows an almost perfectly linear scaling of PYVALE’s memory usage as a function of image size for a single pair of reference and deformed images. The subset-size and subset-step were fixed at 21 and 10 pixels respectively. These quantities do have an impact on memory consumption, especially as subset-steps approach a single pixel. For the subset-step and subset-size used here, by far the most dominant quantity that affects memory usage is the image size. It is also worth noting here that thread count has little affect on memory consumption for standard




Fig. 9 (Top) Absolute runtime and relative speedup (Bottom) for increasing thread count for the 32000×32000 pixel SEM image pair.

problems, with thread local memory requirements mainly being dominated by subset information and optimization parameters, which are small compared to the shared image and interpolant buffers. 50 GB of memory was required in total.

5 Comparison of Runtime to Other Open-source Software Packages

PYVALE was compared against three other open-source, local-subset-based DIC codes: OpenCorr [13], DICE [45], and Ncorr [12]. All comparisons shown in this section were performed on the same workstation described in the previous section (AMD Ryzen Threadripper 7980X with 64 cores, 256 GB of RAM, running Ubuntu 24.04.2 LTS x86_64). OpenCorr and DICE are implemented in C++, while Ncorr is a MATLAB extension that is heavily dependent on interactions within a GUI. There is some limited scripting functionality, but this is generally restricted to loading reference, deformed images and ROI selection. In contrast, both OpenCorr and DICE support extensive scripting and can be executed independently of a GUI, enabling a


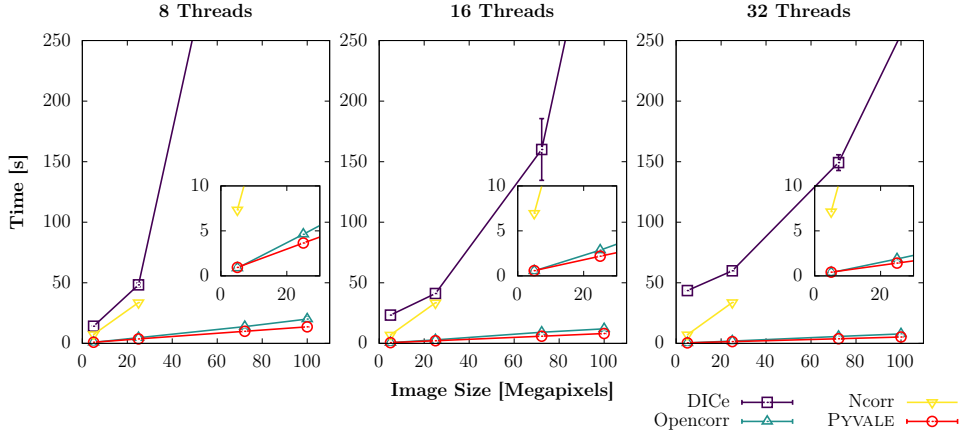


Fig. 10 Peak memory consumption of PYVALE for varying image sizes. Thread count was fixed at 64.

systematic study of runtime across varying image sizes and thread counts. Although the DIC codes examined here differ in documentation quality and maintenance status, each was installed and configured according to its respective documentation to ensure proper operation. Input parameters were aligned based on the available documentation wherever possible to facilitate a fair, like-for-like comparison. All scripts and results are openly available through [46].

Two comparisons were conducted; The first used a synthetically generated speckle pattern stretched radially outward from the image center with a maximum extension of 15 pixels along each edge. Image dimensions of $2,000 \times 2,500$, 5000×5000 , $8,500 \times 8,500$, 10000×10000 pixels were used. Attempts with 32000×32000 were made, however all codes, except PYVALE, crashed due to exceeding the workstations amount of allocated RAM. The second comparison looked at performance under less ideal conditions involving larger displacements. Images of the same sizes to those listed above were generated from centered crops of the SEM dataset used in the previous section. A rectangular ROI was used with excluded boundaries of 100, 200, 300, and 400 pixels for the respective image dimensions.

Synthetic Speckle Pattern



SEM Images

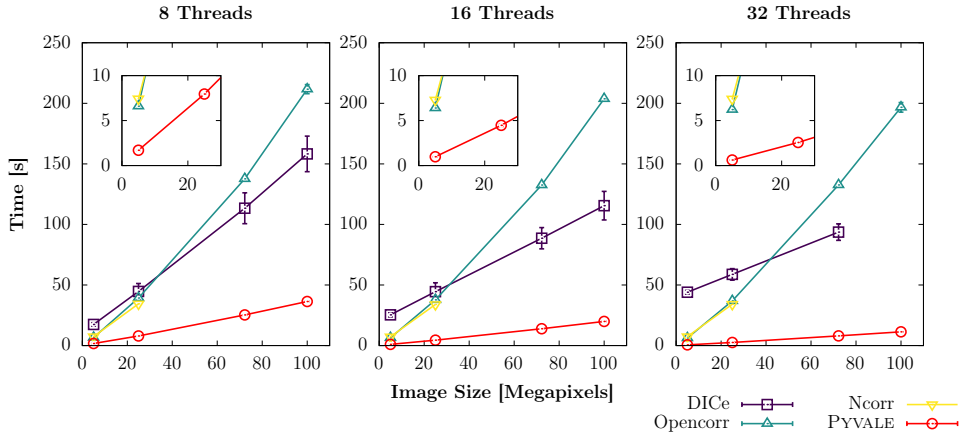


Fig. 11 Total end-to-end runtime the two image sets used in the open-source code comparison. The top row shows results for the synthetically generated speckle pattern for varying image sizes and thread count. Bottom row shows runtime for the SEM images used in previous sections. Missing data points are due to programs failing due to exceeding RAM limit.

Correlation parameters were standardized across all codes where possible. The maximum subset optimization iteration count was set to 40, and the cutoff precision set to 0.01. The subset-size for PYVALE, OpenCorr, and DICe was fixed at 31 with a subset-step of 15. Ncorr, which uses circular subsets, employed a radius of 15 pixels (for a 30 pixel diameter). Thread counts of 8, 16, and 32 were tested. All C++ codes were compiled using GCC with the ‘-O3’ optimization flag. Ncorr was executed using a standard MATLAB 2025b installation.

The Ncorr MATLAB extension requires a seed location for the RG-DIC process in

each thread; these were selected manually within the GUI with best efforts made to distribute them evenly across the ROI. OpenCorr’s initial displacement estimates were computed using a Scale-Invariant Feature Transform (SIFT) approach. Although OpenCorr also supports FFTCC initialization, its window size matches the subset-size used in the non-linear optimization, preventing accurate detection of displacements larger than the subset in a single pass. Implementing a multi-window FFTCC algorithm within the C++ setup script used by OpenCorr would have been possible but deemed to be beyond the expected scope for a typical user. Consequently, SIFT initialization was used for the large-deformation SEM images, while FFTCC initialization was applied to the synthetic speckle patterns which have smaller deformations. For DICE, a feature-matching initialization was used for both image sets.

Presented runtimes are for full end-to-end execution performed for 3 repetitions for each code. For Ncorr, reported times exclude any GUI interactions and include only initialization and correlation. For DICE, OpenCorr, and PYVALE, runtimes are from script execution to termination. This naturally includes an initialization, correlation, terminal and file Input/Output (I/O). Runtime comparisons for the three thread counts on both the synthetic speckle and SEM datasets are shown in Fig. 11. Missing data points indicate cases where a code has crashed across all three repeats due to running out of available memory. PYVALE and OpenCorr show comparable performance for the synthetic speckle images with runtime decreasing as thread count increases. In these tests, both codes use an FFTCC on a single subset sized (31×31) window to estimate rigid displacements before performing non-linear optimization to compute displacement fields. Ncorr’s runtime remains nearly constant across all thread counts. The long initialization time (relative to the correlation time) does not change significantly with increasing thread count, in contrast to the correlation time, which decreases as more threads are used. A breakdown into initialization and correlation runtimes is provided across Fig. 12 and Fig. 13.

Another noticeable difference between datasets is the increased runtime of OpenCorr for SEM images relative to synthetic patterns. This is attributed to the initialization approach: their FFTCC initialization efficiently captures small deformations (within a single subset) in synthetic images, whereas the SIFT-based approach required for large SEM deformations introduces much greater computational overhead compared to the multi-window FFT approach used in PYVALE.

Along with runtime performance, installation time and setup complexity are also important considerations for typical DIC users. For PYVALE, the goal is to make the installation process as straightforward as possible. By a considerable margin, Python remains the most widely used programming language for engineering and scientific applications [47, 48]. As long as the user has a compatible version of Python and pip installed, PYVALE can be installed with a single command: `pip install pyvale`. This eliminates the need to manually manage commonly used C++ dependencies for image processing, such as OpenCV [49] and FFTW [50], or work with build systems such as CMake [51].

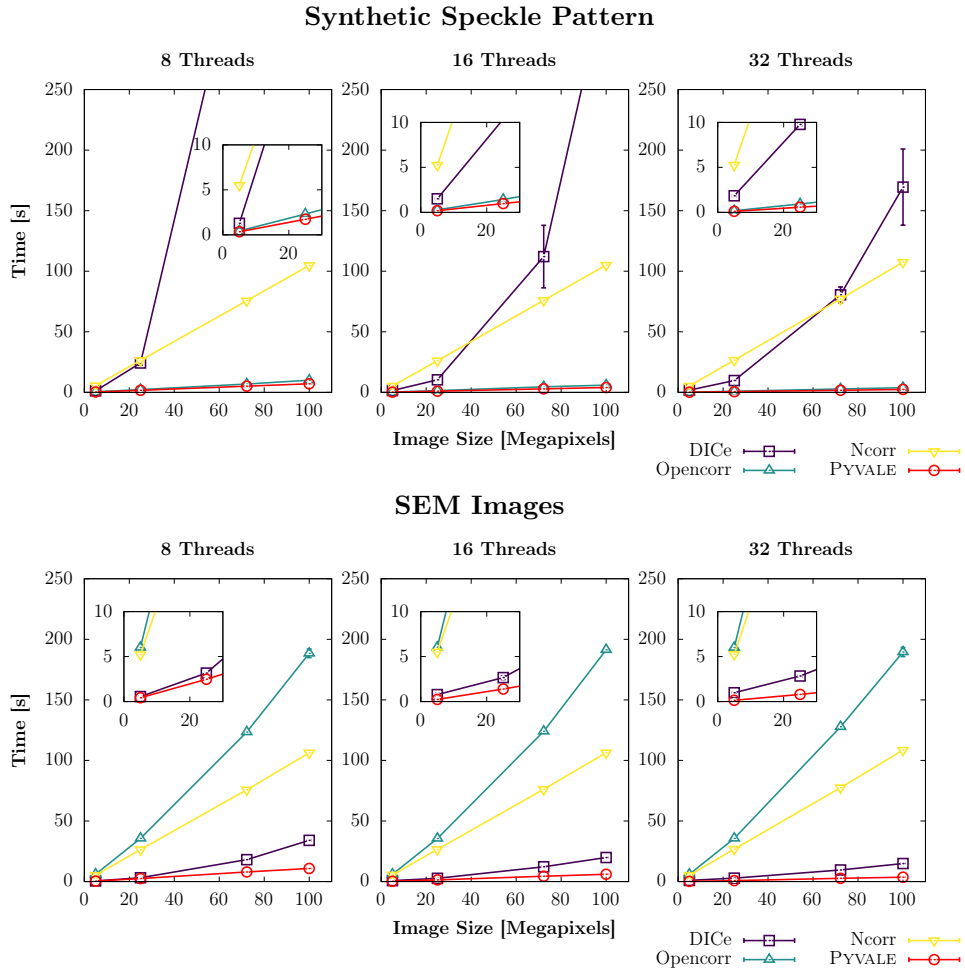
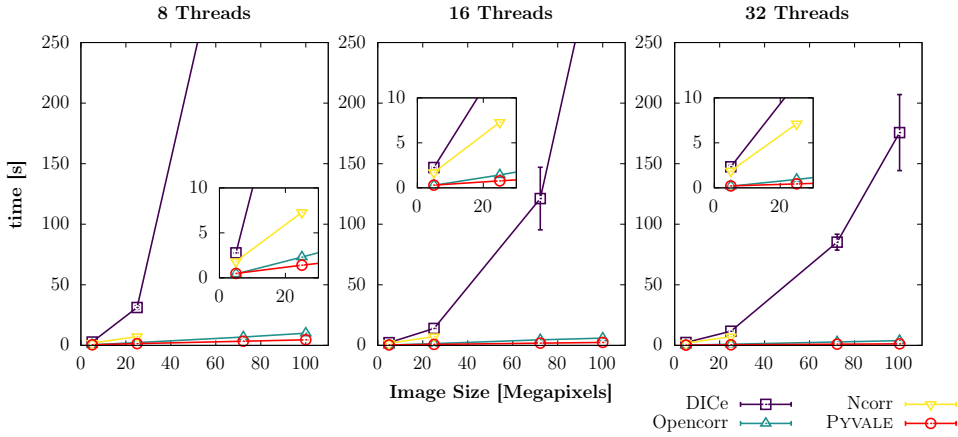


Fig. 12 Initialization time for the two image sets used in the open-source code comparison. The top row shows results for for the synthetically generated speckle pattern with a hydrostatic strain for varying image sizes and thread count. Bottom row shows runtime for the SEM images used in previous sections. Missing data points are due to programs failing due to exceeding RAM limit. PYVALE and OpenCorr perform with similar speeds for the synthetic speckle pattern as both are using a FFTCC cross correlation approach. DICe’s feature matching performs much better for the SEM images, likely due to the artifacts in the material providing suitable features to correlate between reference and deformed images. Ncorr shows consistent times across thread counts, it is therefore likely that the section of code responsible for the initialization does not make use of multi-threading.

6 Future Development

Specifically for the 2D DIC module, several enhancements are planned. One is the integration of an automated reference-image updating scheme [52], in which the reference image evolves incrementally through the sequence of deformed images to accommodate large strains. Another planned improvement is an automated system

Synthetic Speckle Pattern



SEM Images

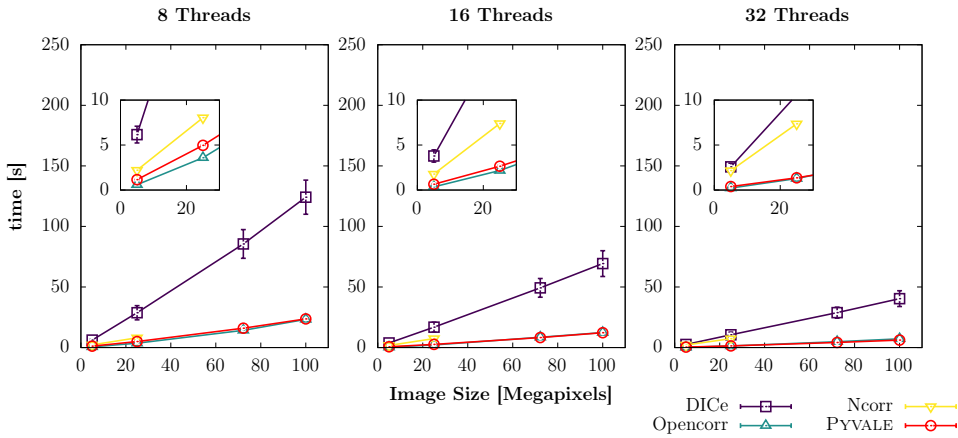


Fig. 13 Correlation time for the two image sets used in the open-source code comparison. The top row shows results for the synthetically generated speckle pattern with a hydrostatic strain for varying image sizes and thread count. Bottom row shows runtime for the SEM images used in previous sections. Missing data points are due to programs failing due to exceeding RAM limit. Ncorr, OpenCorr and PYVALE all have similar correlation times and show a decrease in runtime for increasing thread count. DICe was shown to be considerably slower in for both image sets used here.

for partitioning large images into smaller sequential chunks that can be loaded into memory one at a time during processing. This will allow PYVALE to handle arbitrarily large images efficiently, even on machines with typical amounts of memory (e.g., 8–32 GB of RAM). While PYVALE’s DIC engine is currently limited to 2D, a key milestone on the development roadmap is the implementation of stereo DIC with full stereo calibration capabilities. Another major objective is to extend PYVALE to benefit from GPUs, especially to accelerate the multi-window FFT initialization utilizing the availability of specific FFT libraries [53–55] for this hardware. Several existing DIC

codes have already demonstrated substantial performance gains using this type of hardware [56–58]. Alongside continued development of the DIC engine itself, efforts will also focus on integrating the engine into the broader planned PYVALE ecosystem. This includes coupling the DIC engine with synthetic image deformation driven by FEA simulations, as well as enabling experimental design investigations through optimization of sensor placement and lighting using ray-tracing and rasterization graphics engines.

Beyond the planned new features listed above, ongoing work will focus on improving usability of the existing code base. Requests for new features, examples, or bug reports can be submitted through the PYVALE GitHub repository [59].

7 Summary

PYVALE has been designed to be an open-source and easy-to-use package for DIC with a simple installation procedure. By creating a simple Python interface that hooks into faster lower level languages for computationally heavy tasks, PYVALE is developing into a tool that is both accessible to new users and capable of handling the increasingly large images and datasets without compromising on either scalability or speed. Here, we have demonstrated the capabilities of PYVALE through: (i) a numerical validation against the DIC challenge 2.0, (ii) the application of PYVALE to gigapixel resolution SEM-DIC images, with displacements approaching 1000’s of pixels. (iii) a performance comparison to other local-subset based open source codes demonstrating PYVALE is the most performant code for large SEM-DIC images.

While metrological performance is comparable for other local-subset DIC codes, what sets PYVALE (and global DIC codes such as muDIC [15]) apart is the free open-source, easy-to-use Python interface. This, coupled with strong numerical performance will hopefully allow PYVALE to become a highly utilized code within the experimental mechanics community.

PYVALE is a free and open-source package, and contributions from the DIC community are encouraged, including for commercial use (PYVALE is distributed under the MIT license). Users are welcome to explore, modify, and extend the software to suit their own research needs, and to deploy as many copies of PYVALE as they see fit, whether on individual workstations and laptops for small-scale testing or on high-performance computing clusters for systematic sweeps of DIC parameters. Looking ahead, we aim for this project to support ongoing community-driven development and to provide a solid platform for future improvements in open-source DIC.

8 Authorship Contribution Statement

- **Joel Hirst:** Methodology, software writing, software validation, Article writing - original draft, Article writing - review & editing.
- **Lloyd Fletcher:** Conceptualization, software writing, software validation, article writing - review & editing, project administration, funding acquisition.
- **Adel Tayeb:** Conceptualization, software validation, project administration.

- **Lorna Sibson, Megan Sampson, Wiera Bielajewa:** Software writing, software validation.
- **Ben Poole:** Article writing - original draft, Experimental SEM work, software validation.
- **Michael Atkinson, Alex Marsh, Rory Spencer, Rob Hamill:** Software validation,
- **Cory Hamlin, Allan Harte:** Funding acquisition, project administration.

9 Acknowledgments

We acknowledge support from UKRI through the Future Leaders Fellowship scheme, grant MR/Y015916/1. We also acknowledge support from the UK Engineering and Physical Sciences Research Council (EPSRC), grant EP/W006839/1.

10 Data Availability Statement

The source code for PYVALE can be accessed through the Github repository [59]. All scripts and results for the comparisons made in this work can be found in [46].

References

- [1] Schreier, H., Orteu, J.-J., Sutton, M.A.: Image correlation for shape, motion and deformation measurements: Basic concepts, theory and applications. Image Correlation for Shape, Motion and Deformation Measurements (2009)
- [2] Pan, B.: Digital image correlation for surface deformation measurement: Historical developments, recent advances and future goals. Institute of Physics Publishing (2018). <https://doi.org/10.1088/1361-6501/aac55b>
- [3] Jones, E.M.C., Iadicola, M.A. (eds.): A Good Practices Guide for Digital Image Correlation. International Digital Image Correlation Society, ??? (2018). <https://doi.org/10.32720/idics/gpg.ed1>
- [4] Peters, W., Ranson, W.: Digital imaging techniques in experimental stress analysis. *Optical engineering* **21**(3), 427–431 (1982)
- [5] Chu, T., Ranson, W., Sutton, M.A.: Applications of digital-image-correlation techniques to experimental mechanics. *Experimental mechanics* **25**(3), 232–244 (1985)
- [6] Chen, D.-J., Chiang, F.-P., Tan, Y., Don, H.: Digital speckle-displacement measurement using a complex spectrum method. *Applied optics* **32**(11), 1839–1849 (1993)
- [7] Pan, B.: Reliability-guided digital image correlation for image deformation measurement. *Appl. Opt.* **48**(8), 1535–1542 (2009) <https://doi.org/10.1364/AO.48.001535>
- [8] Schreier, H.W., Braasch, J.R., Sutton, M.A.: Systematic errors in digital image correlation caused by intensity interpolation. *Optical engineering* **39**(11), 2915–2921 (2000)
- [9] MatchID: MatchID Software. <https://www.matchid.eu> (2025)
- [10] GmbH, L.: DaVis: Integrated imaging acquisition and processing software. <https://www.lavision.de/en/products/davis-software/>. Version 11 (2025)
- [11] Correlated Solutions, Inc.: Correlated Solutions. Accessed: 2025-10-27 (2025). <https://www.correlatedsolutions.com/>
- [12] Blaber, J., Adair, B., Antoniou, A.: Ncorr: Open-source 2d digital image correlation matlab software. *Experimental Mechanics* **55**(6), 1105–1122 (2015) <https://doi.org/10.1007/s11340-015-0009-1>
- [13] Jiang, Z.: Opencorr: An open source library for research and development of digital image correlation. *Optics and Lasers in Engineering* **165**, 107566 (2023)

- [14] Turner, D.Z.: DICe: Digital Image Correlation Engine. GitHub repository, <https://github.com/dicengine/dice>. Version 2.0; accessed on 19 June 2025 (2020)
- [15] Olufsen, S.N., Andersen, M.E., Fagerholt, E.: μ dic: An open-source toolkit for digital image correlation. *SoftwareX* **11**, 100391 (2020)
- [16] Lockwood, W., Reynolds, A.P.: Use and verification of digital image correlation for automated 3-d surface characterization in the scanning electron microscope. *Materials Characterization* **42**(2-3), 123–134 (1999)
- [17] Kammers, A.D., Daly, S.: Digital image correlation under scanning electron microscopy: methodology and validation. *Experimental Mechanics* **53**(9), 1743–1761 (2013)
- [18] Lindfeldt, E., Ekh, M., Cvetskovski, K., Schilke, M.: Using dic to identify microscale strain fields from in-situ sem images of a pearlitic steel. *Experimental mechanics* **54**(9), 1503–1513 (2014)
- [19] Liu, B., Lan, S., Li, J., Fang, Q., Ren, Y., He, W., Xie, H.: Digital image correlation in extreme conditions. *Thin-Walled Structures* **205**, 112589 (2024)
- [20] Kammers, A., Daly, S.: Small-scale patterning methods for digital image correlation under scanning electron microscopy. *Measurement Science and Technology* **22**(12), 125501 (2011)
- [21] Jin, H., Lu, W.-Y., Korellis, J.: Micro-scale deformation measurement using the digital image correlation technique and scanning electron microscope imaging. *The Journal of Strain Analysis for Engineering Design* **43**(8), 719–728 (2008)
- [22] Mello, A., Book, T., Nicolas, A., Otto, S., Gilpin, C., Sangid, M.: Distortion correction protocol for digital image correlation after scanning electron microscopy: emphasis on long duration and ex-situ experiments. *Experimental Mechanics* **57**(9), 1395–1409 (2017)
- [23] Poole, B., Marsh, A., Lunt, D., Hardie, C., Gorley, M., Hamelin, C., Harte, A.: High-resolution strain mapping in a thermionic lab6 scanning electron microscope. *Strain* **60**(5), 12472 (2024) <https://doi.org/10.1111/str.12472>
<https://onlinelibrary.wiley.com/doi/pdf/10.1111/str.12472>
- [24] Computer-Aided Validation Laboratory Team, U.: Pyvale: The Python Validation Engine. <https://computer-aided-validation-laboratory.github.io/pyvale/index.html>. Release 2025.8.1, MIT License (2025)
- [25] Hunter, J.D.: Matplotlib: A 2d graphics environment. *Computing in Science & Engineering* **9**(3), 90–95 (2007) <https://doi.org/10.1109/MCSE.2007.55>
- [26] Inc., P.T.: Collaborative data science. Plotly Technologies Inc., Montreal, QC

- (2015). <https://plot.ly>
- [27] Scarano, F., Riethmuller, M.L.: Iterative multigrid approach in piv image processing with discrete window offset. *Experiments in Fluids* **26**(6), 513–523 (1999)
- [28] Scarano, F., Riethmuller, M.L.: Advances in iterative multigrid piv image processing. *Experiments in fluids* **29**(Suppl 1), 051–060 (2000)
- [29] Liberzon, A., Lasagna, D., Aubert, M., Bachant, P., Käufer, T., jakirkham, Bauer, A., Vodenicharski, B., Dallas, C., Borg, J., tomerast, ranleu: OpenPIV/openpiv-python: OpenPIV - Python (v0.22.2) with a new extended search PIV grid option. Zenodo (2020). <https://doi.org/10.5281/zenodo.3930343>
- [30] Lunt, D., Thomas, R., Roy, M., Duff, J., Atkinson, M., Frankel, P., Preuss, M., Da Fonseca, J.Q.: Comparison of sub-grain scale digital image correlation calculated using commercial and open-source software packages. *Materials Characterization* **163**, 110271 (2020)
- [31] Reinecke, M.: PocketFFT: A Fast Fourier Transform Library. GitHub. Accessed: 2025-09-19 (2018). <https://github.com/mreineck/pocketfft>
- [32] Brigham, E.O.: *The Fast Fourier Transform and Its Applications*. Prentice-Hall, Inc., ??? (1988)
- [33] Marquardt, D.W.: An algorithm for least-squares estimation of nonlinear parameters. *Journal of the society for Industrial and Applied Mathematics* **11**(2), 431–441 (1963)
- [34] Levenberg, K.: A method for the solution of certain non-linear problems in least squares. *Quarterly of applied mathematics* **2**(2), 164–168 (1944)
- [35] Pan, B., Xie, H., Wang, Z., Qian, K., Wang, Z.: Study on subset size selection in digital image correlation for speckle patterns. *Optics express* **16**(10), 7037–7048 (2008)
- [36] Dagum, L., Menon, R.: Openmp: an industry standard api for shared-memory programming. *IEEE computational science and engineering* **5**(1), 46–55 (1998)
- [37] Reu, P.L., Toussaint, E., Jones, E., Bruck, H.A., Iadicola, M., Balcaen, R., Turner, D.Z., Siebert, T., Lava, P., Simonsen, M.: Dic challenge: Developing images and guidelines for evaluating accuracy and resolution of 2d analyses. *Experimental Mechanics* **58**(7), 1067–1099 (2018) <https://doi.org/10.1007/s11340-017-0349-0>
- [38] Reu, P.L., Blaysat, B., Andò, E., Bhattacharya, K., Couture, C., Couty, V., Deb, D., Fayad, S.S., Iadicola, M.A., Jaminion, S., *et al.*: Dic challenge 2.0: developing images and guidelines for evaluating accuracy and resolution of 2d analyses: focus

- on the metrological efficiency indicator. *Experimental Mechanics* **62**(4), 639–654 (2022)
- [39] Helm, J., Bossuyt, S., Reu, P., Turner, D., Luan, L., Lava, P., Siebert, T., Simonsen, M.: Stereo-dic challenge 1.0—rigid body motion of a complex shape. *Experimental Mechanics* **64**(7), 1073–1106 (2024)
- [40] Croom, B.P., Burden, D., Jin, H., Vonk, N., Hoefnagels, J., Smaniotto, B., Hild, F., Quintana, E., Sun, Q., Nie, X., *et al.*: Interlaboratory study of digital volume correlation error due to x-ray computed tomography equipment and scan parameters: an update from the dvc challenge. *Experimental Mechanics* **61**(2), 395–410 (2021)
- [41] Grédiac, M., Blaysat, B., Sur, F.: A critical comparison of some metrological parameters characterizing local digital image correlation and grid method. *Experimental Mechanics* **57**(6), 871–903 (2017)
- [42] Blaysat, B., Neggers, J., Grediac, M., Sur, F.: Towards criteria characterizing the metrological performance of full-field measurement techniques: application to the comparison between local and global versions of dic. *Experimental Mechanics* **60**(3), 393–407 (2020)
- [43] Zukić, D., Jackson, M., Dimiduk, D., Donegan, S., Groeber, M., McCormick, M.: Itkmontage: a software module for image stitching. *Integrating Materials and Manufacturing Innovation* **10**(1), 115–124 (2021)
- [44] Works, S.R.: Mean Absolute Deviation. Springer (2008). https://doi.org/10.1007/978-0-387-32833-1_250
- [45] Turner, D.Z.: Digital image correlation engine (dice) reference manual. Sandia Report SAND2015-10606 O, Sandia National Laboratories (2015)
- [46] Hirst, J.: Dataset for Pyvale: A Fast, Scalable 2D Digital Image Correlation (DIC) Engine Capable of Handling Large-Scale Images. <https://doi.org/10.5281/zenodo.17670411>
- [47] Cass, S.: The top programming languages 2025. *IEEE Spectrum* (2025). Accessed: 2025-11-19
- [48] Jansen, P.: TIOBE Programming Community Index. <https://www.tiobe.com/tiobe-index/>. Accessed: 2025-11-19 (2025)
- [49] Bradski, G.: The OpenCV Library. *Dr. Dobb’s Journal of Software Tools* (2000)
- [50] Frigo, M., Johnson, S.G.: The design and implementation of fftw3. *Proceedings of the IEEE* **93**(2), 216–231 (2005)
- [51] Kitware, Inc. and Contributors: CMake Reference Documentation. Kitware,

- Inc., (2025). Kitware, Inc.. Accessed: 2025-10-16. <https://cmake.org/cmake/help/latest/index.html>
- [52] Bing, P., W, D., Xia, Y.: Incremental calculation for large deformation measurement using reliability-guided digital image correlation. In: Optics and Lasers in Engineering, vol. 50, pp. 586–592 (2012). <https://doi.org/10.1016/j.optlaseng.2011.05.005>
- [53] Corporation, N.: cuFFT. GPU-accelerated Fast Fourier Transform (FFT) library (2025). <https://developer.nvidia.com/cufft>
- [54] Advanced Micro Devices, I.: AOCL-FFTW (Fastest Fourier Transform in the West) (2025). <https://www.amd.com/en/developer/aocl/fftw.html>
- [55] Tolmachev, D.: Vkfft-a performant, cross-platform and open-source gpu fft library. IEEE Access **11**, 12039–12058 (2023)
- [56] Couty, V., Witz, J.-F., Lecomte-Grosbras, P., Berthe, J., Deletombe, E., Brieu, M.: Gpucorrel: a gpu accelerated digital image correlation software written in python. SoftwareX **16**, 100815 (2021)
- [57] Blug, A., Regina, D.J., Eckmann, S., Senn, M., Bertz, A., Carl, D., Eberl, C.: Real-time gpu-based digital image correlation sensor for marker-free strain-controlled fatigue testing. Applied Sciences **9**(10), 2025 (2019)
- [58] Zhang, L., Wang, T., Jiang, Z., Kemao, Q., Liu, Y., Liu, Z., Tang, L., Dong, S.: High accuracy digital image correlation powered by gpu-based parallel computing. Optics and Lasers in Engineering **69**, 7–12 (2015)
- [59] Computer-Aided Validation Laboratory: pyvale. GitHub repository (2025). <https://github.com/Computer-Aided-Validation-Laboratory/pyvale>