

INFLUENCE OF CHARGE SHAPE ON THE HETEROGENEOUS MECHANICAL RESPONSE OF CARBON FIBRE SHEET MOULDING COMPOUNDS

Henri Schwalm^{1,2} and Pascal Hubert^{1,2}

¹ Structures and Composite Materials Laboratory, McGill University, Montréal, Canada

² Centre de recherche sur les systèmes polymères et composites à haute performance (CREPEC)

Keywords: Sheet Moulding Compound, Digital Image Correlation, Compression Moulding

ABSTRACT

Long discontinuous fibre reinforced Carbon Fibre Sheet Moulding Compounds (CF-SMC) are quick and economical to process into finished parts, they are stiff and fairly strong, but their heterogeneous morphology results in high variability in mechanical performance across a single part and from part to part. Since the shape of the charge has a great influence on the in-mould flow, and thus the part morphology, its design is a valuable tool that has yet to be fully understood. To best evaluate its effect, this work introduces a high-resolution multi-view Digital Image Correlation technique which allows a closer look at the heterogeneity in CF-SMC. It emphasizes measurement fidelity and maps the DIC data onto image data of the virgin part to analyze all samples in a common context. The technique is applied to coupon tensile tests using samples cut from a flat plate moulded at different amounts of charge coverage. Meso-scale heterogeneity remains high across all charge coverages with a slight increase in areas with long flow paths. On the macro-scale, averaging DIC strain fields helps diminish the apparent stiffness variability of CF-SMC.

1 INTRODUCTION

Carbon Fibre Sheet Moulding Compounds (CF-SMC) occupy a highly auspicious sweet spot between manufacturability and performance due to their material architecture based on randomly oriented strands (ROS). While discontinuous nature of ROS allows for their use in efficient compression moulding processes, their long, bundled fibres transfer loads effectively, giving them a stiffness comparable to equivalent continuous fibre quasi-isotropic laminates [1]. Through the use of high fibre volume fractions of carbon fibres, CF-SMC are distinctly different from conventional SMCs and achieve a mechanical performance that makes them promising for structural applications. Compared to their continuous fibre counterparts, however, CF-SMCs suffer from inferior strength and high variability in stiffness and strength. This unpredictable apparent brittleness can be attributed to the heterogeneous morphology of the material itself. Digital Image Correlation (DIC) used on both sides of a tensile test specimen, yields full field strain measurements featuring strain concentrations up to 10 times the magnitude of the global, even if the opposing side shows no such peak [2]. They are particularly common in regions, where the stress concentrations on the ends of strands aligned with the loading direction coincide with low-stiffness regions such as transversely oriented strands. Such a coalescence of critical features has been shown to likely be the initiation point for final failure in the matrix-dominated failure modes ROS are known for [2], [3].

Next to the material architecture itself, the flow compression behaviour during compression moulding has the most significant impact on said variability. Long flow paths due to a low coverage of the mould by the initial charge have been shown to contribute to a reorientation of the fibre bundles, causing anisotropic mechanical behaviour [4], and an increasing occurrence of internal defects and detrimental morphology features, such as fibre swirling and kinking [3]. For these advanced SMC variants, a low-flow charge configuration is hence often preferred [1], [3]. However, for intricate parts, this requires complex shapes to be cut from a roll of precursor material and assembled into a charge much resembling a preform. The resulting material offcuts and increased labour go against the quick and economical zero-waste principles paramount to the motivation behind SMC. While the macro-scale influence of flow compaction on anisotropic mechanical performance is reasonably well understood, the detailed results

of DIC have yet to be put into context with the morphology created through differing degrees of in-mould flow. Consequently, this work investigates how different charge coverage configurations influence the heterogeneous strain fields measured during mechanical testing of the moulded parts as a step toward balancing low performance variability and feasibly simple charge design. A novel DIC-based method is introduced, which aims to maximize the amount of usable data, its fidelity and accurately correlates it back to the original structure. It's applicability to CF-SMC is demonstrated experimentally for samples with different degrees of initial charge coverage.

2 METHOD DEVELOPMENT: MULTI-VIEW HIGH-RESOLUTION DIGITAL IMAGE CORRELATION

Due to the highly heterogeneous mechanical response of ROS, DIC has emerged as the most insightful strain measurement technique for this material architecture. To tell the complete story, from charge design to part morphology to locally varied mechanical performance, this work puts particular emphasis on image resolution, and how said resolution can effectively contribute to a more accurate and detailed data set. While the core principles of the method may be applied to a plethora of mechanical tests, tensile tests according to ASTM D3039 are the use case for this example.

2.1 Hardware Configurations and Calibration

Two different configurations of the same set of hardware were used to utilize their respective advantages. They are depicted in Figure 1.

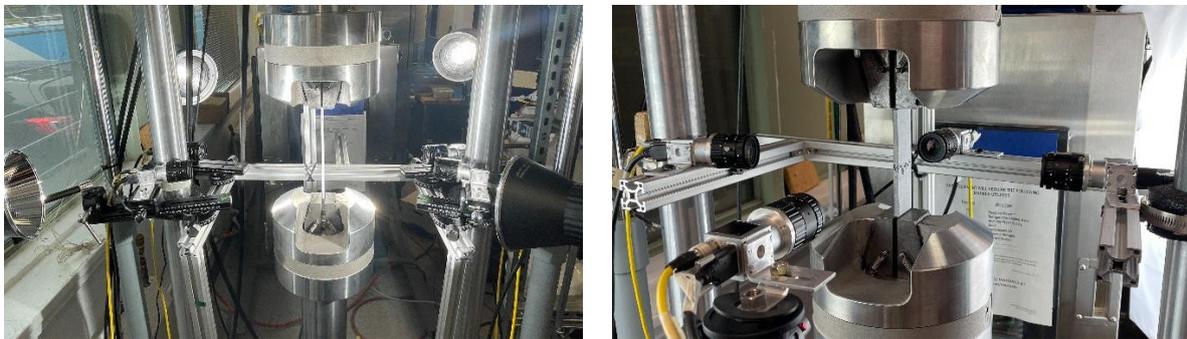


Figure 1: Digital Image Correlation setup. a) Dual 3D DIC. b) Quadruple 2D DIC.

Firstly, four 20 MP machine vision cameras grouped into two stereo camera pairs are arranged on both sides of the tensile specimen (cf. Figure 1 a), much like in the work of Johanson et al. [2]. This way, both opposing main surfaces of the sample can be measured using 3-dimensional (3D) DIC, which yields out-of-plane displacement data in addition to the in-plane one. While this is usually not of great importance in tensile tests, it can be useful to initially troubleshoot slight grip misalignment, and to spot surface deformations due to cracks, as shown in Figure 2 a), and b), respectively. Furthermore, using 3D DIC over 2D DIC based on a single camera is generally considered best practice, as the stereo camera calibration process inherently compensates for lens distortions [5]. A double-sided calibration target is used to introduce the data into a shared coordinate system.

However, once validated against 3D DIC, and a calibrated extensometer, the 2D DIC approach was found to be similarly accurate, and thus it is the foundation for the second hardware configuration shown in Figure 1 b). Here, the same four cameras are arranged perpendicularly to all four sides of the specimen, enabling a close look at through-thickness strains as they appear on the cut side surfaces of the sample, as shown in Figure 3. Calibration is also simplified, as each camera inherently has shared reference points with its neighbour, such as the ends of the speckled region. Assuming the surfaces as perfectly perpendicular, a single length measurement within one image allows the data to be mapped into 3D space, e.g. to validate structural simulations.

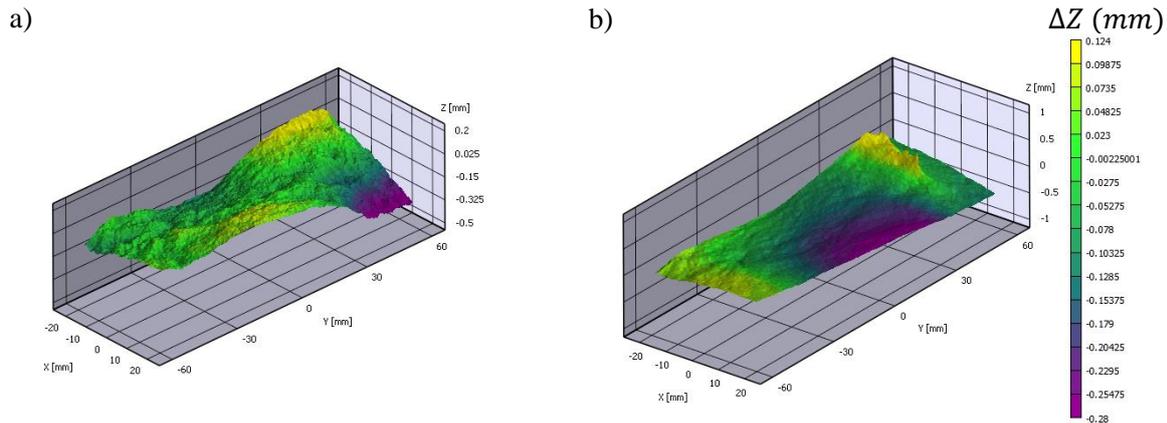


Figure 2: Vertical offset from the neutral plane on the surface of a CF-SMC ASTM D3039 tensile test sample. a) Minute grip misalignment in torsion made visible by DIC, not apparent to the naked eye. b) Microcrack bulges outward halfway through the test.

Synchronizing the mechanical load data with the image data is usually done by the DIC acquisition software, which samples an analogue feed from the load frame through a data acquisition device every time it acquires an image set. As most of these software packages are not meant for the use with more than two cameras, the given approach instead relies on the mechanical testing system for timing. Every time the load frame system stores a data point, it physically triggers all cameras through a digital output connected to the general-purpose input-output pins that can be found on most machine vision cameras, separate from the data connection.

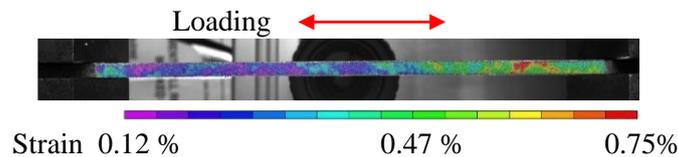


Figure 3: Highly heterogeneous strain in through-thickness direction of a CF-SMC ASTM D3039 tensile specimen visible on its side.

2.2 On the Use of Rolling Shutter Cameras

Modern digital machine vision cameras are the current standard for data acquisition in DIC. For low-speed applications, they can be directly connected to a computer through standardized USB3 Vision or GigE Vision interfaces, they feature monochrome image sensors to avoid introducing measurement biases from Bayer Filter arrays used in color cameras, and they do not feature mechanical shutters for exposure timing. Instead, their electronic shutter simply queries the sensor photoreceptors for the desired amount of time. Two main types of electronic shutters can be found inherent to the respective image sensors. Global shutters trigger all pixels of the sensor simultaneously and are generally considered best practice for DIC [5]. However, rolling shutter cameras are far more common and generally much lower priced than their equivalent global shutter counterparts. Furthermore, they typically achieve lower noise levels at the same imaging conditions [6]. The exposure is read out row by row of pixels along the sensor height with a slight delay from one to the next. As such, one side of the image depicts a slightly later point in time than the other.

To fit four high resolution cameras within the project budget, and since their lower noise levels contribute to the desired high data fidelity, the developed method uses rolling shutter cameras featuring Sony IMX183 sensors. To assure that one image truthfully depicts one single time point of the mechanical test, the setup and test parameters were adapted accordingly. When the expected displacement within the image frame is balanced against the rolling shutter delay, the error can be minimized to the point of it being negligible. In the case of the ASTM D3039 tensile test, the standard

allows for a large leeway in testing speed, allowing for such an adaptation. For rolling shutter DIC, the maximum crosshead speed is chosen such that the maximum displacement expected within the DIC area of interest (Aoi) remains smaller than 0.01 pixels on the camera sensor during the time the shutter needs to scan the Aoi. The measure of 0.01 pixels is common confidence goal in DIC. The orientation of the cameras is thus of great importance, since there is no noticeable lag within a single row of pixels. By arranging the cameras in portrait orientation, the scanning direction is transverse to the tensile one, keeping scanning length of the Aoi small, since the samples are less wide than they are long. For the configuration used in this setup the overall rolling shutter delay and the resulting maximum test speed are 40 ms and 1 mm/s respectively.

2.3 Local Modulus and Denoising Low Strain Results

Predicting the exact failure location of ROS samples based on low-strain DIC results has been shown to be challenging [2], since some fatal strain concentrations develop later during the test. Nonetheless, high strain regions within the linear elastic region can be valuable indicators of the low local performance due to unfavourable morphology. Tang et al. [7] introduce the nominal modulus for ROS, averaging the strain values belonging to one cross-section of the sample, calculating a modulus using the global load within the linear elastic region, and repeating this for the length of the gauge section. Similarly, this work uses the local effective modulus of every single subset as primary stiffness performance indicator using the global stress to calculate. This neglects the inhomogeneous true stress distribution within a cross-section but is mainly meant to denoise the strain data in a physically relevant way; while it is common to apply spatial domain filtering when calculating full field DIC to even out errors introduced by image noise, each image is usually treated independently from one another. This way, the time domain goes mostly unused. Smoothing each value along the time axis, or, more appropriately, the load axis, provides further denoising based on the principle of linear elasticity. Combining a low amount of spatial filtering with linear regression along the load preserves the surface heterogeneity detail while cleaning up noisy strain fields. The inverse of the slope of that regression is the local effective modulus, herein simply called local modulus. The regression is performed between 1000 and 3000 microstrain, as in ASTM D3039, but may prove promising for even lower strains to identify problem regions in the morphology in non-destructive tests or in-use component tests.

2.4 Data Matching Using Feature Matching and Homography

The surface appearance of ROS already contains a plethora of information about the morphology of the specimen that can be obtained without resorting to more laborious and expensive techniques such as Micro Computed Tomography. Johanson et al. [2] found critical surface morphology features such as strand ends along the loading direction closely coinciding with DIC strain concentrations. For more complex surface appearances, such connections may not always be as simple to draw. Building on the idea, the proposed method systematically and accurately correlates surface image data with DIC data for use in further analysis and predictive methods. The highly automated method computes a series of transformation matrices and is shown in Figure 4. Aside from the DIC test images, a number of photographs must be obtained during sample preparation, all of which should be taken perpendicular to the sample with minimal distortion. Cross-polarization of the light sources and lenses helps control unwanted reflections and glare. Firstly, high resolution pictures of all surfaces of the uncut part serve as a reference onto which the data is projected. After cutting, the tensile specimens are photographed from both sides before and after applying the DIC speckle pattern in the exact same position using a fixed reference. The DIC strain data is exported from the analysis software as a point cloud with x- and y coordinates as they pertain to the image pixels of the first DIC test image. Through use of the Scipy Python library, a convex hull around this point cloud is computed, which constitutes the perimeter of the part of that image that is within the Aoi and used as a mask to isolate the contents (Step 1). The necessary perspective transformation to map this image fragment onto the speckled flat reference image is then calculated using Feature Matching and Homography algorithms from the open-source computer vision library OpenCV. Unique keypoints in both images are identified in both images using Scale Invariant Feature Transform (SIFT). Their compatibility is then assessed by a nearest neighbour search based on algorithms designated Fast Library for Approximate Nearest Neighbours (FLANN).

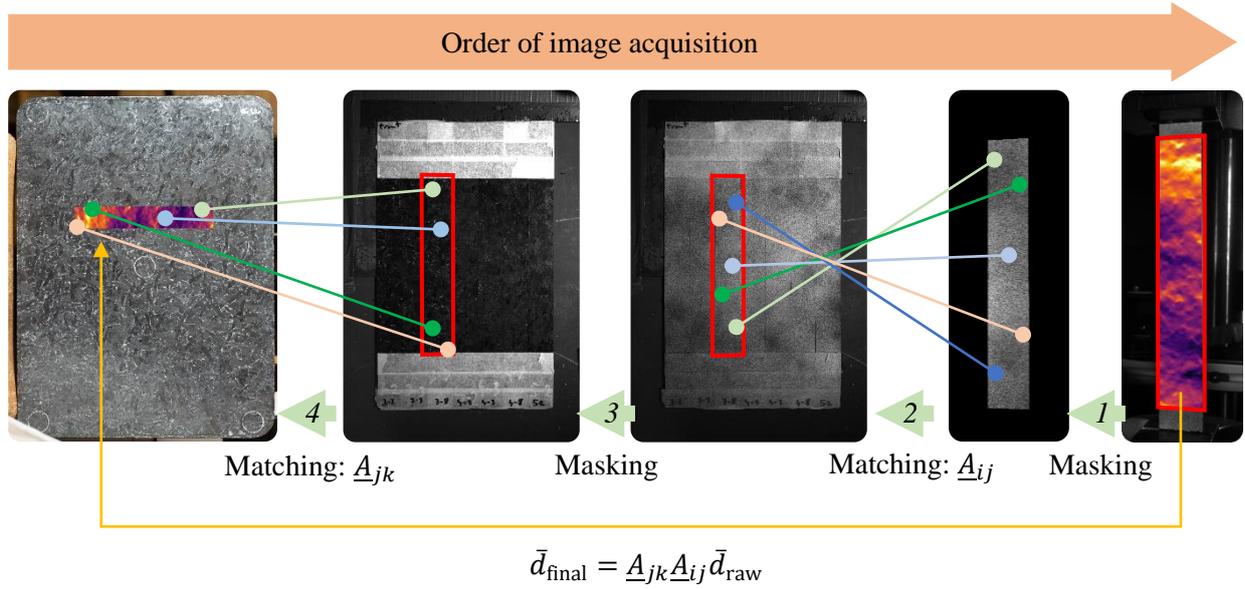


Figure 4: Working principle of DIC data association algorithm. Homography yields Transformation matrices \underline{A} to transform data set \bar{d} into the coordinate system of the original sample image. The multi-colored lines represent matching keypoints in an image pair.

Given the coordinates of the sets of matching keypoints, the homography algorithm outputs a 3 x 3 transformation matrix linking the images together (Step 2). Transforming the original convex hull onto the speckled flat reference image using that matrix, it can be once again used as a mask (Step 3). Since the un-speckled version of the that image has the sample in the exact same location, the matrix effectively maps the DIC data onto the surface image data. To see this in the context of the entire sample and make statements about the local mechanical performance relative to the charge design, the Aoi, as extracted from the un-speckled flat reference image is then projected onto the un-cut plate using the same feature matching and homography algorithms (Step 4). Through this method, all data points are now in a shared coordinate system pertaining to the image of the uncut part, can be grouped into sections of the part, and evaluated according to this grouping.

3 MATERIALS AND METHODS

A commercial CF-SMC with a nominal fibre fraction of 50 % by weight and a vinyl-ester matrix was used in this study. Fibre strands were 25.4 mm long and about 3 mm wide before compression. The material was processed in a hardened steel, chrome plated mould for a flat plate with dimensions of 305 mm by 406 mm inside of a 250-ton hydraulic press. Three different charge configurations were investigated, each covering the entire length of the long side of the mould, and covering 100 %, 75 %, and 50 % of the short side of the mould, respectively. To keep the final plate thickness equal at 3.6 mm, these charges were made up of 3, 4, and 6 layers of SMC each, respectively. They were placed toward one edge of the mould to promote significant material flow representative of half of a symmetric larger ply that would likely be centrally placed in an industrial application. According to supplier instructions and previous resin characterization, one plate for each charge coverage were moulded at a temperature of 135 °C at 100 bars of pressure for 5 minutes each. Using a wet cutting diamond saw, five rectangular tensile samples were cut from each plate in accordance with ASTM D3039 at 25.4 mm wide. All samples were cut along the material flow direction. The samples were tensile tested using the method detailed in Section 2, with two samples per configuration tested using the dual 3D DIC configuration, and three samples using the quadruple 2D DIC configuration. The local modulus data from the latter three was then mapped onto images of the plates and evaluated statistically.

4 RESULTS AND DISCUSSION

The results of the tensile tests are presented below with special attention given to the diversity of data contained within multi-view high resolution DIC. The amount of usable information yielded by a single test is weighed against the standard test using an extensometer for strain measurement. The statistical capabilities of the method are demonstrated by evaluating different sections of a part separately from each other across multiple samples.

4.1 Information Contained Within a Single Test

The high heterogeneity in ROS is most often mainly attributed to its fibre orientation distribution, which is inherently random, but affected by flow compaction. This suggests that the stiffness inconsistency is most prominent on the mesoscale (close to the strand length) and should even out as one “zooms out” toward the macro scale. As the length of standard rectangular tensile coupons is multiple times that of the strand length, the distribution of strain values within an entire gauge section should yield much more information than a single strain gauge, which essentially just samples one very particular region of the sample. Figure 5 shows the distribution of roughly 100,000 strain values from the subsets within the linear elastic range of a single test, and how they relate to a range of multiple repetitions of a test with the same charge coverage where strain is measured by one strain gauge each.

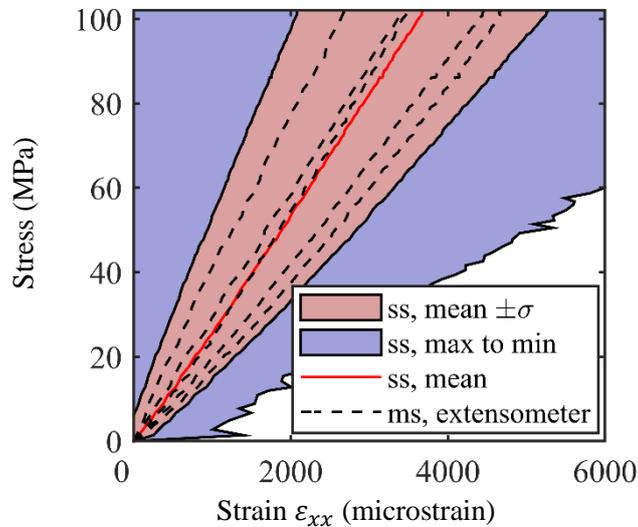


Figure 5: Stress-strain curve comparing the DIC statistics of a single sample (ss) against multiple samples (ms) measured with an extensometer. Red and blue colored envelopes encompass all values contained within the range of standard deviation σ and across the whole gauge section, respectively.

The modulus obtained from the mean of all subsets in the single specimen closely approximates the mean of the modulus value from the multiple tests, at 37 GPa and 36 GPa, respectively. The standard deviation strain envelope around that mean encompasses all individual measurements and could thus potentially serve as a conservative estimate of the standard deviation to be expected from a test set using extensometers. While the least strained subsets are not of great technical interest, the strain concentration regions may indicate potential early failure sites, giving importance to the minimum local modulus. This limited example shows that examining both the distribution of all strain values, and the different distributions in multiple samples or sample regions is of interest and can reveal previously unseen trends.

4.1 Heterogeneity of Tensile Elastic Properties

Figure 6 shows a visual representation of both specimen sides from all charge configurations, the local modulus data precisely mapped back onto the plate surface images. Some surface features, like the ends of longitudinal strands cause strain concentrations and hence low stiffness regions, as observed by other authors.

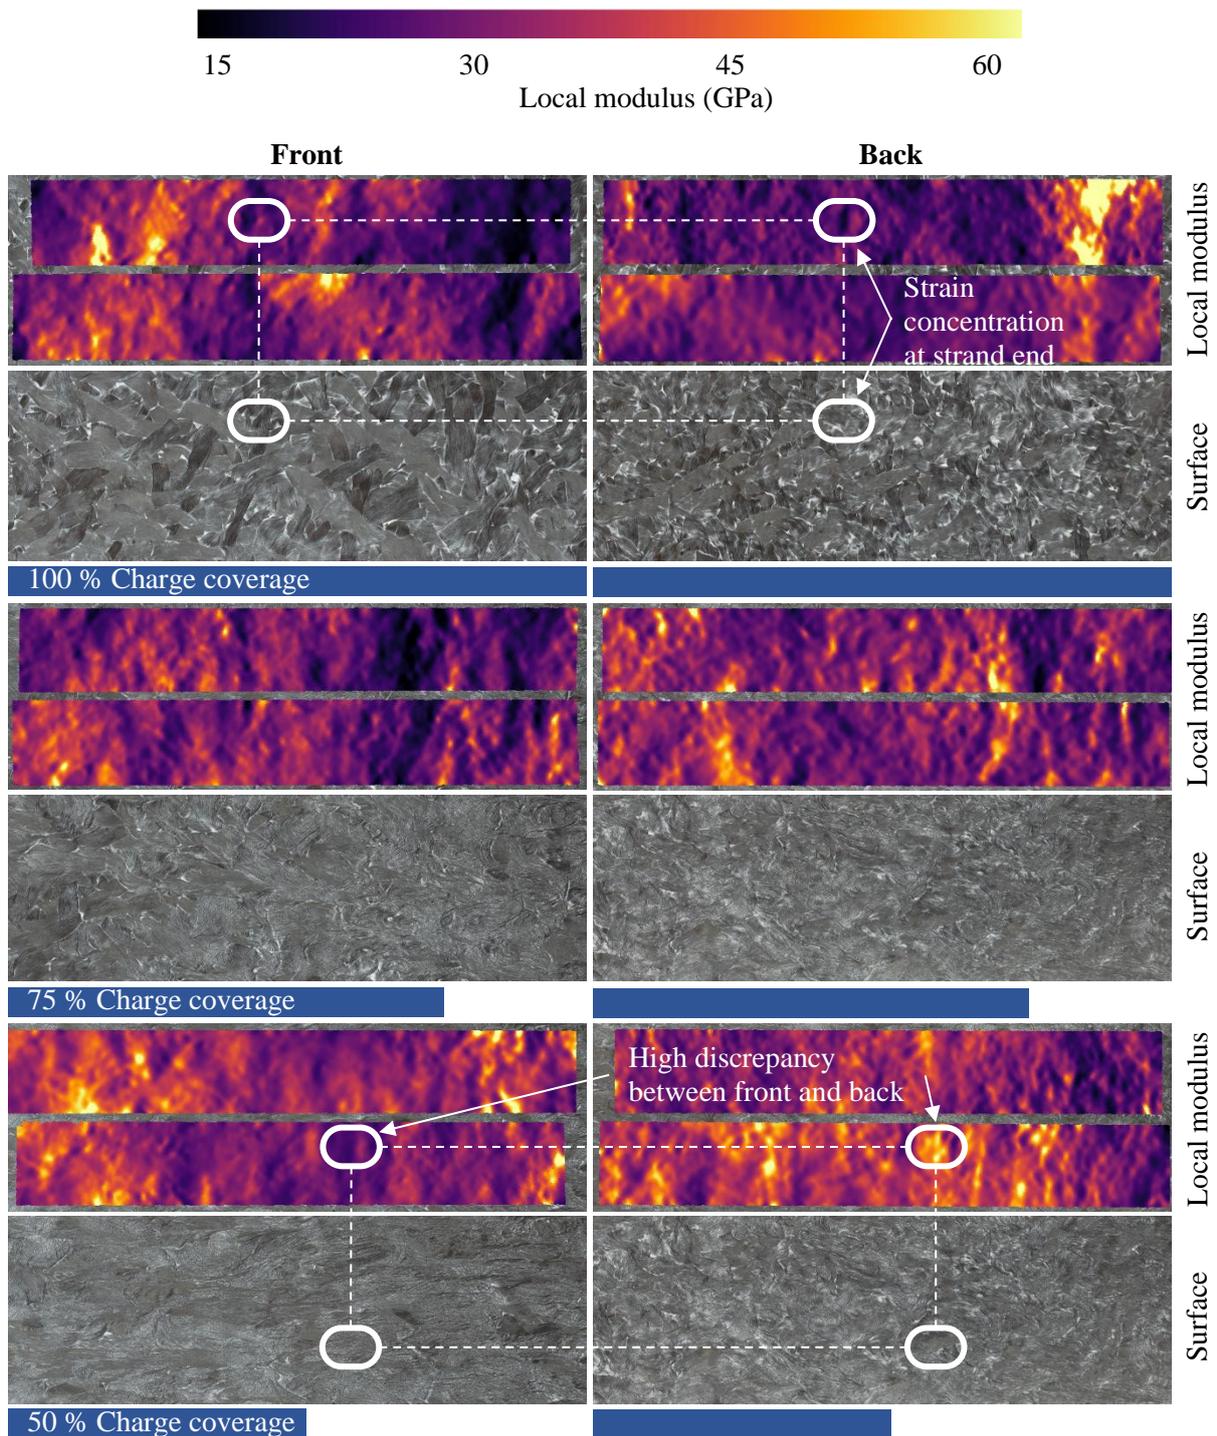


Figure 6: DIC data mapped onto surface image data. Section that was initially covered by the charge is marked in dark blue.

Other such concentrations are harder to explain from looking at the surface morphology, especially in high flow specimen with less easily distinguishable features. Their cause might be found in a more complex interaction of multiple strands as well as morphology features hidden beneath the surface. Furthermore, large differences in strain between the front and back of a sample are visible, as previously reported in literature [2], [7]. The distribution of all strain values from all samples of the different charge coverages are presented in Figure 7 a).

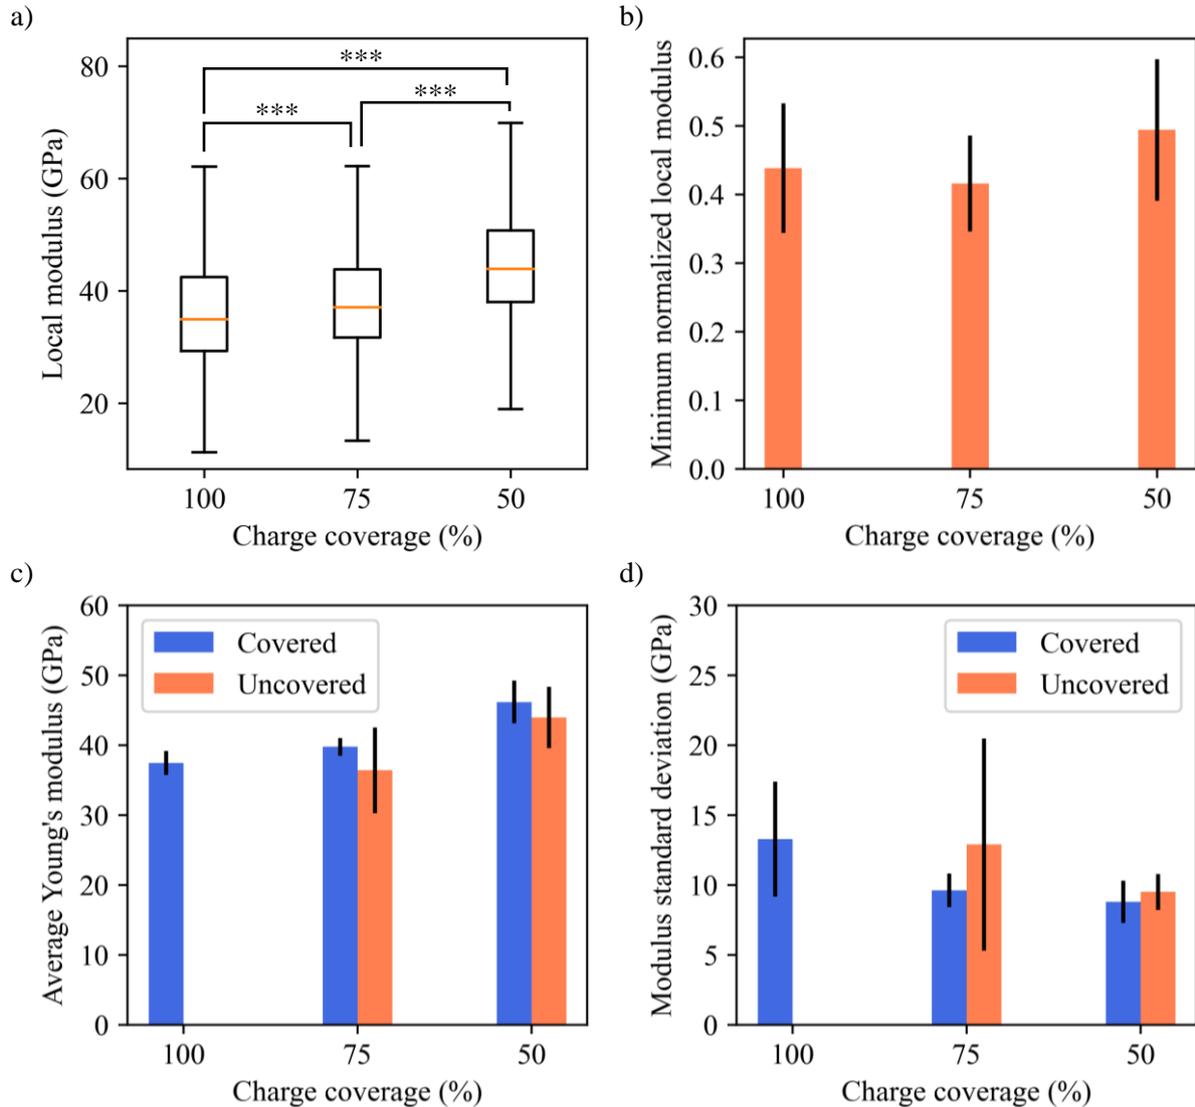


Figure 7: Statistical evaluation of DIC results. a) Local modulus distribution of all subsets from all samples. b) Minimum moduli per sample normalized with overall modulus. c) Mean of subsets within different sections of the sample. d) Standard deviation of local modulus within different sections.

*** denotes $p < 0.001$.

The anisotropy induced by the in-mould flow is made apparent by the overall higher stiffness values at lower charge coverages. The overall heterogeneity however is largely unaffected, with similar quartile ranges across all charge coverages. A very slight skewedness toward lower stiffness values is visible in all three distributions. Local low stiffness regions may be indicative of strain concentrations which tend to act as failure initiation sites. Thus, Figure 7 b) compares the minimum local stiffness values for each individual sample normalized by its overall stiffness. Similar values are seen for all charge configurations. The 50 % charge coverage samples exhibit slightly less severe concentrations than the other two, suggesting that the increased defects in high-flow samples previously reported in literature do not seem to outweigh the impact of the inherent morphology randomness. Figure 7 c) and d) evaluate each single tensile sample individually. The sample is divided into a section of the mould that was initially covered by the charge, and one that was initially uncovered and filled with material through in-mould flow. Figure 7 c) focuses on the mean modulus of each individual section. At 100 % charge coverage there is no initially uncovered section, but the coefficient of variation between samples at 4 % is much lower than is usual for this test in ROS. This supports the theory that most of the apparent stiffness variability is limited to the mesoscale, and low-flow samples can behave quite predictably on the

macro-scale. Initially uncovered sample sections seem to exhibit a slightly lower mean modulus and higher heterogeneity, as demonstrated by the standard deviations of local modulus within each of the individual sections (cf. Figure 7 d).

9 SUMMARY AND CONCLUSIONS

A multi-view high resolution DIC specifically for the investigation of the heterogeneous mechanical response of CF-SMC and other ROS material architectures was introduced and tested. It aims to maximize the amount and quality of useable data extracted from mechanical tests through its data fidelity and reconfigurability. A feature matching and homography algorithm was introduced to precisely associate this data with image data of the original part. This facilitates the design and evaluation of techniques that predict mechanical performance based on morphology data. Furthermore, it allows the data of multiple tests to be analyzed in a shared coordinate system and grouped based on the shape of the part and its charge. Less accurate charges predictably result in higher stiffness along the flow directions, but notably also slightly elevated heterogeneity in parts of the sample with long flow paths. The simple metric of charge coverage does hence not seem sufficient as the sole indicator for process induced anisotropy and heterogeneity of ROS. Instead, the present work emphasizes the local meso-scale heterogeneity over the macro-scale one, a phenomenon which has been mostly attributed to local fibre orientation constellations [2], [3]. The tailored design of charges seems a valuable tool to affect these orientation distributions through in-mould flow. In conjunction with numerical process modelling to predict fibre the orientation evolution, this DIC data matching technique can be applied to many more charge configurations, sample orientations, and mechanical tests to facilitate a methodical CF-SMC charge design that makes the most of this material class.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the financial support provided by the Natural Sciences and Engineering Research Council of Canada as well as PRIMA Québec. We kindly thank our industrial partner for their continued support and guidance.

REFERENCES

- [1] S. B. Visweswaraiah, M. Selezneva, L. Lessard, and P. Hubert, "Mechanical characterisation and modelling of randomly oriented strand architecture and their hybrids – A general review," *J. Reinf. Plast. Compos.*, vol. 37, no. 8, pp. 548–580, 2018, doi: 10.1177/0731684418754360.
- [2] K. Johanson, L. T. Harper, M. S. Johnson, and N. A. Warrior, "Heterogeneity of discontinuous carbon fibre composites: Damage initiation captured by Digital Image Correlation," *Compos. Part A Appl. Sci. Manuf.*, vol. 68, pp. 304–312, 2015, doi: 10.1016/j.compositesa.2014.10.014.
- [3] L. M. Martulli, L. Muyschondt, M. Kerschbaum, S. Pimenta, S. V Lomov, and Y. Swolfs, "Carbon fibre sheet moulding compounds with high in-mould flow: Linking morphology to tensile and compressive properties," *Compos. Part A Appl. Sci. Manuf.*, vol. 126, 2019, doi: 10.1016/j.compositesa.2019.105600.
- [4] A. D. Evans, C. C. Qian, T. A. Turner, L. T. Harper, and N. A. Warrior, "Flow characteristics of carbon fibre moulding compounds," *Compos. Part A Appl. Sci. Manuf.*, vol. 90, pp. 1–12, Nov. 2016, doi: 10.1016/j.compositesa.2016.06.020.
- [5] International Digital Image Correlation Society, E. M. C. Jones, and M. A. Iadicola, "A Good Practices Guide for Digital Image Correlation," p. 94, 2018, [Online]. Available: <https://doi.org/10.32720/idics/gpg.ed1/print.format>
- [6] Teledyne Photometrics, "Rolling Shutter vs. Global Shutter." 2014. [Online]. Available: <https://www.photometrics.com/learn/advanced-imaging/rolling-vs-global-shutter>
- [7] H. Tang *et al.*, "Correlation between failure and local material property in chopped carbon fiber chip-reinforced sheet molding compound composites under tensile load," *Polym. Compos.*, vol. 40, no. S2, pp. E962–E974, 2019, doi: 10.1002/pc.24767.