strain WILEY

The particle image tracking technique: An accurate optical method for measuring individual kinematics of rigid particles

Vincent Richefeu [ORCID] | Gaël Combe

Univ. Grenoble Alpes, CNRS, Grenoble INP, 3SR, Grenoble, F-38000, France

Correspondence
Vincent Richefeu, Laboratoire 3SR, Univ. Grenoble Alpes, Grenoble, France.
Email: vincent.richefeu@3sr-grenoble.fr

Abstract

We present a new approach to assess the two-dimensional motion of rigid bodies in granular materials. Although it was adapted from digital image correlation technique, the heart of the presented technique relies on specific treatments related to the discrete nature of grain-displacement fields. The code called TRACKER has been developed to process the digital images and measure the in-plane displacement and rotation of each individual grain from one image to another. A remarkable feature is the use of a specific strategy that allows tracking all particles, without *losing* any of them (which is a typical problem when tracking assemblies of discrete particles over many images). This is achieved by a two-step procedure, where, in case of problematic tracking of a grain, the size of the search zone is increased in an adaptive manner, that is, taking into account the results of tracking in the neighbourhood of the particle. The accuracy of the measured displacements and rotations was tested on both perfect synthetic images and digital photographs of a sheared assembly of grains. An automatic procedure that corrects the lens distortion further improves the quality of the measurements. The accurate assessment of the grain kinematics opens very interesting perspectives, especially in the study of displacement fluctuations in granular media.

KEYWORDS

accuracy check, digitial image correlation, discrete materials, kinematic assessment, subpixel

1 | INTRODUCTION

The apparatus named $1\gamma2\varepsilon$ [1], developed at Laboratoire 3SR in Grenoble, France, is a two-dimensional (2D) parallelepiped that allows for any quasi-static straining on 2D granular analogue materials—called Schneebeli materials [2]—Figure 1. The rod packing stands vertically and undergoes the gravity field. It has been used to study, for instance, the structure evolution of granular materials under complex loading paths [3] or the effect of particle shapes on the macroscopic strength [4]. One special feature of this device is that the front face of the specimen is exposed. Pictures can thus be shot during loading tests and then analysed to extract full 2D kinematics of individual grains. In the 90s, this extraction was performed by means of 12×9 cm silver photographies in which particle translations and rotations were manually measured using a stereo comparator device. The displacement (translation) accuracy was quite good
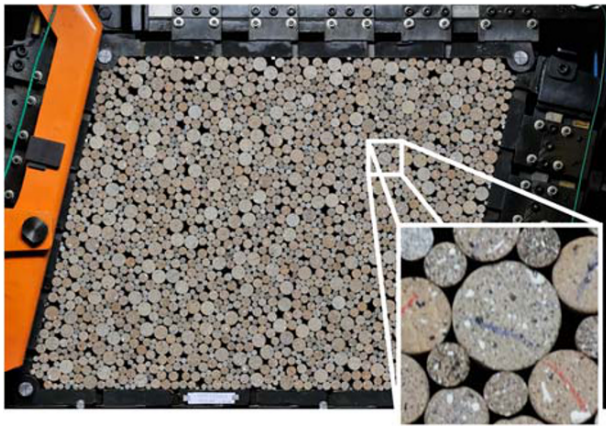
FIGURE 1   Assembly of about 2,000 wood rods placed vertically in the device $1\gamma2\varepsilon$ device[1]. For this example, the initial size of the packing 55×60 cm, and the mean diameter of the rods is $\sim$13 mm

(in the order of 3×10^{-3} mean diameter), but the time required to treat each photograph was prohibitive (typically 1.5 h per photo to manually click on 1,000 rods). At the start of this century, analogue cameras were replaced by digital ones. Thus, a new numerical photogrammetry technique to obtain semi-automated measurements of kinematics was used[5]. The use of digital cameras made the shots easier (no need of the photography silver film processing step); however, the image analysis procedure was not able to reduce the process time. Between each couple of successive pictures, even if most particles *move* almost like continuum mechanics dictates, many of them have unpredictable displacements, and it was necessary to manually "help" the process to find the right position and rotation of grains. Moreover, the accuracy was very coarse (in the order of the pixel, which correspond to 10^{-2} of particle's mean diameter with the 16 MPixels digital camera used at this time).

It was therefore decided to develop a new way to assess full 2D kinematic from digital photographs of granular samples made of Schneebeli rods with the following specifications: fully automated measurement of grain kinematic (translations and rotations) able to handle more than 1,000 digital photographs and above all, a displacement accuracy strongly increased compared with that obtained with the previous method[5]. The technique presented can be regarded as the *Lagrangian* version of the particle image velocimetry technique that aims to determine from pairs of successive images the velocity of tracer particles seeded within a mass-flow of particles, with an *Eulerian* manner. With this respect, we called it the particle image tracking (PIT).

It is important to emphasise that it is not so straightforward to track particles in an individual and Lagrangian way. In fact, known codes, although efficient and accurate in measuring continuous fields, fail to successfully perform this task of particle tracking. This is mainly due to the unpredictable nature of the movement of some particles as illustrated in Figure 2; a specificity of the plastic deformation of dense (hyperstatic) granular media where each particle is subjected to strong geometric constraints by its neighbourhood. In addition to the description of the PIT method below, other related aspects have been appended so as not to disturb the reading. These are (a) the quality of the pattern signals, (b) the algorithm parallelisation, and (c) the correction of distortion and parallax of acquired images; these topics are addressed in Appendices B1, C1, and D1, respectively.
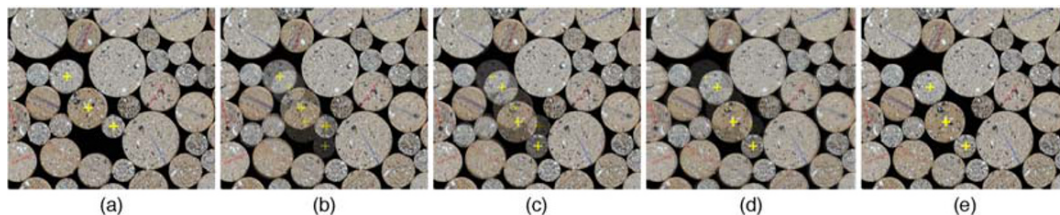

FIGURE 2   (a–d)Sequence of images showing the "jump" of three grains—marked with yellow crosses—within a bidimensional granular assembly (Schneebeli rods). These jumps are characterised by very large displacements in comparison with the average displacement of the neighbouring strains. Such events, which have a more or less pronounced amplitude, are not uncommon during a macroscopic deformation and constitute what could be qualified as fluctuation of displacements

2 | THE PIT TECHNIQUE

Over the last few years, there has been substantial development in quantitative imaging analysis. One of them is the *digital image correlation* (DIC), initially suggested by Chu et al.[6]. It has become a powerful technique that provides reliable kinematic measurement fields, for example, displacement, acceleration, strain, and strain-rate fields. This technique is a fully nonintrusive measurement tool (provided that image acquisition is itself nonintrusive) that can be used to follow the straining of patterns of a wide range of continuous materials such as metals, polymers, ceramics, and concretes. The particular case of granular materials (fine sands, silts, etc.) can be treated as a continuum when the followed pattern includes several grains[7]. In this peculiar case, the resulting displacements and their gradients have to be regarded as averaged transformations over the tracked pattern.

Another usage of the DIC can be envisaged; it consists to track the grain's displacements and rotations, one by one, within the material. However, problems arise when the DIC is applied at this grain scale: the motion of individual grain is not regular because of their mutual geometric exclusion, and it does not strictly mirror the kinematics imposed at the packing boundaries[8,9].For this specific purpose, we have developed an image processing software called TRACKER where the speckles on a digital image of each grain are used to track their displacements u, v and orientation θ from their original position r and orientation θ.

2.1 | The DIC technique customised for rigid motions

The basic concept of the DIC relies on finding the best similarity between the subsets of two images. The subset from the first image is transformed (e.g., displaced, stretched, tilted, and/or locally opened) according to a set of parameters $\{\mathcal{P}\}$ in order to be correlated with the second image. The best correspondence is tested trough a cross-correlation coefficient φ. It can be the *zero mean normalised cross correlation* (ZNCC) coefficient, which has the benefit of not being sensitive to changes in amplitude and shift of the correlated signal—it indeed means that any change in lighting in the photographs will not affect the cross correlation. This coefficient reads

$$\varphi(\{P\}) = \frac{\sum_{\mathcal{S}}(I_1(\boldsymbol{x}_1) - \bar{I}_1)(I_2(\boldsymbol{x}_2) - \bar{I}_2)}{\sqrt{\sum_{\mathcal{S}}(I_1(\boldsymbol{x}_1) - \bar{I}_1)^2 \sum_{\mathcal{S}}(I_2(\boldsymbol{x}_2) - \bar{I}_2)^2}}, \tag{1}$$

where I_n is the grey level of the pixel of coordinates $\boldsymbol{x}_n = (x_n, y_n)^T$ of the picture n and $\bar{I}_n$ is the mean grey level computed over the corresponding pixel subsets $\mathcal{S}_n$.

A first specificity of the PIT is to define a subset $\mathcal{S}$ as relative positions instead of a simple grid. In this way, the *patch* can be given any shape that will be used to follow a pattern throughout the rigid movements. It is thus possible to adapt the shape of the patch to the shape of the grains being tracked (Figure 3) and thus optimise the correlated signal.

Let us now focus on the transformation of the vector $\Delta \boldsymbol{x}_1 = \boldsymbol{x}_1 - \boldsymbol{r}$ towards the vector $\Delta \boldsymbol{x}_2 = \boldsymbol{x}_2 - (\boldsymbol{r} + \boldsymbol{u})$, where $\boldsymbol{r}$ is the coordinate of a tracked material point and $\boldsymbol{u}$ is its displacement (see Figure 4). This transformation can be written by means of the transformation gradient operator:

$$\Delta \boldsymbol{x}_2 = \boldsymbol{F} \cdot \Delta \boldsymbol{x}_1. \tag{2}$$

FIGURE 3 Examples of shapes for the tracked pattern defined by a set of pixel positions relative to the particle center. (a) The common rectangular pattern centred on the tracked point, (b) a circular shape that makes possible the definition of a longer optical signal, (c) hollow disk (or ring) that conforms the external part of the particle shape
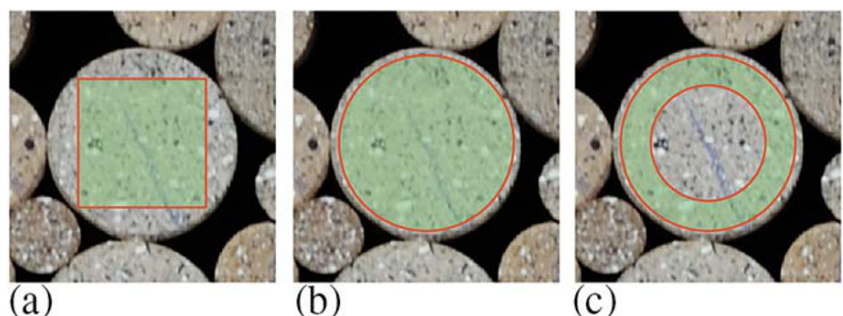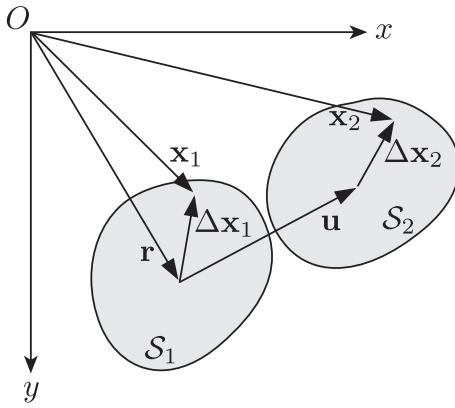
FIGURE 4 Transformation, in Lagrangian description, of a vector Δx_1 at the position r into a vector Δx_2 at the position $r + u$

The transformation gradient expresses as a function of the displacement gradient: $F = \nabla u + \mathbf{1}$, where $\mathbf{1}$ is the identity matrix. Including this latter relation in Equation (2), a position x_1 in the first picture should be moved to the following position in the second picture:

$$x_2 = x_1 + u + \nabla u \cdot \Delta x_1. \tag{3}$$

In this case, the transformation of a pattern is controlled by six parameters held by u and ∇u:

$$\{\mathcal{P}\} = \left[u, v, \frac{\partial u_x}{\partial x}, \frac{\partial u_x}{\partial y}, \frac{\partial u_y}{\partial x}, \frac{\partial u_y}{\partial y} \right]. \tag{4}$$

The tracking procedure in the classical DIC consists in finding for each followed position r, the parameters $\{\mathcal{P}\}$ so that the ZNCC between a pair of signals from two distinct images is maximised. In practice the quantity $1 - \varphi(\{\mathcal{P}\})$ is minimised.

It is interesting to note that Equation (3), which is commonly employed in DIC, might be written otherwise directly from Equation (2), that is, without introducing ∇u. It can thus also simply be written:

$$x_2 = r + u + F \cdot \Delta x_1. \tag{5}$$

In this case, the transformation of a pattern is still controlled by six parameters held by u and F:

$$\{\mathcal{P}\} = \left[u, v, F_{xx}, F_{xy}, F_{yx}, F_{yy} \right]. \tag{6}$$

Given our objective to track the rigid motion of individual grains, it is quite reasonable to state that the transformation gradient is actually a rotation $F \simeq R(\theta)$ with

$$\mathbf{R}(\theta) = \begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix}. \tag{7}$$

The benefit is twofold: first, the assumption of rigid motion is compelled, and second, the number of parameters to be optimised is reduced to three

$$\{\mathcal{P}\} = [u, v, \theta]. \tag{8}$$

If the straining of the pattern needs to be followed (which is not for what the soft TRACKER has been designed for), it is still possible to assess it. In this case, a further optimisation of the displacement u and of its gradient ∇u can be achieved starting with an estimate of the latter:

$$\nabla \boldsymbol{u} \simeq \begin{pmatrix} \cos(\theta) - 1 & -\sin(\theta) \\ -\sin(\theta) & \cos(\theta) - 1 \end{pmatrix}. \tag{9}$$

2.2 | Coarse search

The tracking procedure consists in finding for each grain position $\boldsymbol{r}$, the displacements u, v—with a pixel accuracy—and the rotation θ so that the ZNCC between a pair of signals from two distinct images is maximised. In practice, we minimise the function $1 - \varphi(u, v, \theta)$. To obtain the displacements and rotation with a integer precision, this optimisation is done by testing a set of plausible positions and orientations picked from a rectangular search zone centred on the tracked point that rotates with selected increments. We describe the research zone—that is, the ranges of u and v, but also the angular range—with the set of parameters $\{\Delta u_{\text{left}}, \Delta u_{\text{right}}, \Delta v_{\text{up}}, \Delta v_{\text{down}}, N_{\text{rot}}^-, N_{\text{rot}}^+, \Delta\theta\}$ that are integer values, except $\Delta\theta$. This means that the tested positions range from $x - \Delta u_{\text{left}}$ to $x + \Delta u_{\text{right}}$ horizontally and from $y - \Delta v_{\text{up}}$ to $y + \Delta v_{\text{down}}$ vertically (recall that the y-axis is oriented downwards for digital images); for each tested position, the rotation ranges from $\theta - N_{\text{rot}}^- \Delta\theta$ to $\theta + N_{\text{rot}}^+ \Delta\theta$, with increments of $\Delta\theta$ are tested. All these systematic tests represent $(\Delta u_{\text{left}} + \Delta u_{\text{right}}) \times (\Delta v_{\text{up}} + \Delta v_{\text{down}}) \times (N_{\text{rot}}^- + N_{\text{rot}}^+)$ computations of φ (itself is computed over a more or less large amount of pixels according to Equation 1).

As a result of this methodical step-by-step research, the integer displacement $[u, v]$ and the rotation increment $\Delta\theta$ that induces the highest correlation coefficient φ will be used as a first estimate for the following procedure (subpixel tracking). But, sometimes, due to the granular nature of the sample, a particle can be far from its expected position. This situation can be recognised by a level of correlation that is too low with respect to a given threshold. The strategy implemented in the code TRACKER to solve this problem is the *rescue procedure*. It consists to increase the search zone (including the angular degrees of freedom). A second increase of the search zone is activated when the first rescue attempt is unsuccessful. One can imagine that this procedure slows down the tracking. Fortunately, it is in fact dramatically accelerated when prohibited zones defined by well-found grains are taken into consideration by prohibiting tests on positions screened by the well-found neighbours. Figure 5 pictures this procedure: the rectangular search zone that holds a large number of pixels is reduced by a factor lower than 1/8. In practice, this factor is generally well above.

As a consequence of the rescue procedure, the processing time can be increased. However, in some situations, it is possible to take a small advantage of a certain order in grain processing. For example, if the pattern has a better quality in the larger rods and believing that they are less prone to jumps, the smallest grains can be processed at the end in order to take benefit of the screening of the other grains. Another situation for which the processing order counts is the case where very different displacement amplitudes are expected depending on their position. This is clearly the case for simple shear straining where the displacements linearly increase with the particle height: an order for processing from the bottom to the top of the sample may help to screen the extended search-zone areas following the detection of a rescue. In this way, the first search area can be kept relatively small, and the hope is to reduce, thanks to the screening of well-found particles, the areas that would be expanded. There is therefore a subtle task of finding the right set of parameters.

The choice of the size of the pattern to correlate is also quite critical. Figure 6 illustrates this point. On the left, we see two images shot consecutively as one Schneebeli rod moved more than the others; this grain is marked by a yellow cross. The images on the right show the correlation levels obtained between the original position of the jumping grain

FIGURE 5 Sketch of a rescue procedure. White grains are already well correlated and the red grain needs "rescuing" because of its large displacement between two picture shots. The dashed box is the rectangular search zone that was increased to find the red grain after a first unsuccessful attempt with a very small search zone. Taking into account steric exclusions (light grey zones) due to the local geometry of the assembly of particles and the diameter of the red grain, the blue spaces are places were the centre of the red grain could be. In practice, the blue zone is less the 5% of the area associated to the rectangular search zone
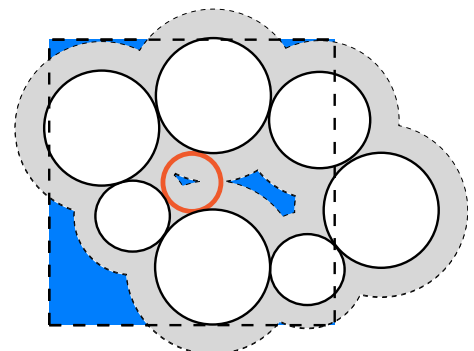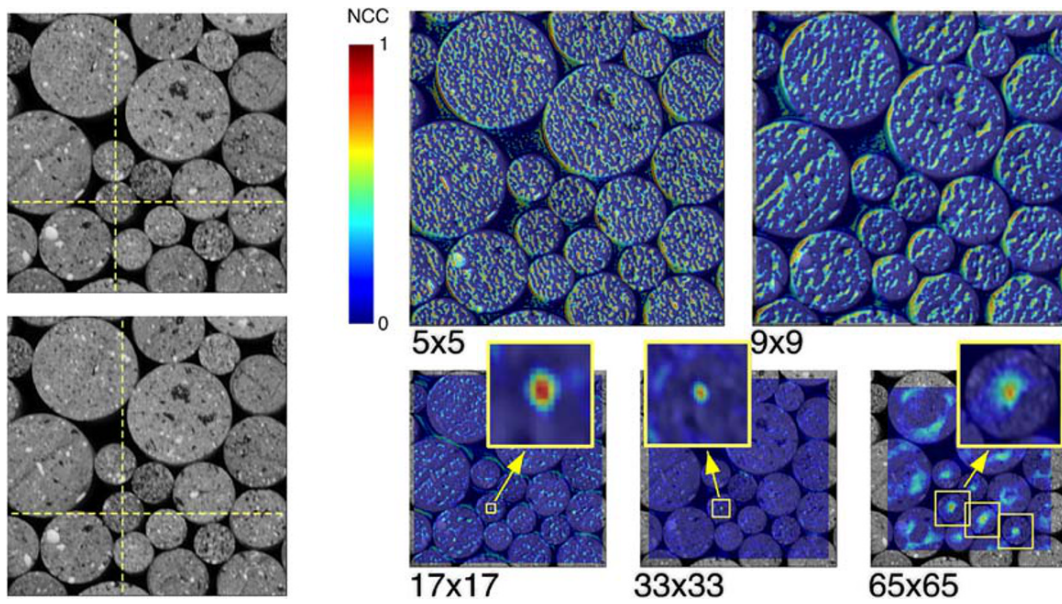
FIGURE 6 Effect of pattern size on the correlation level field. (left) Two images shot at close moments when a particular grain moves strongly with respect to the others. (right) Correlation level colour map (NCC) for different square pattern sizes

in image 1 and all the other achievable pixels in image 2. This is plotted for different sizes of correlated patterns. It can be seen that, up to a size of 17×17 pixels, the correlation-level maps have multiple peaks, which does not clearly and precisely define a good and unique estimate of the new grain position in image 2. For the 17×17 pattern, on the contrary, a single spot is well defined, and it is even more focused for the 33×33 pattern. One might think that the larger the pattern, the better the correlation at the pixel scale, but the map for the 65×65 pattern shows that several spots, centred on other grains can also be found because of a certain sameness of the grains and their overall vicinity. The appropriate size of the patterns is therefore related not only to the speckle but also to the size of the grains. For multidispersed rods in the $1\gamma2\varepsilon$ device, a circular pattern with a slightly smaller radius than the smallest grain should be used so that the tracking accuracy is consistent for all grains.

2.3 | Fine search

To obtain a subpixel resolution, the signal from second image has to be built from interpolated values of grey levels over the subset $\mathcal{S}_2$. In this way, the position $\boldsymbol{x}_2$ in Equations (1) and (3) will be a floating point value whereas it was an integer value in the previous pixel scale procedure. The software TRACKER implements bilinear, bicubic, or biquintic functions to perform the interpolation, partial derivatives being computed by centred finite differences. To find the floating-point values of $[u, v, \theta]$ that best optimise φ, the *Powell*'s conjugate direction method is used[10,11]. This method does not need to define a function for the gradients of φ ($\partial\varphi/\partial u$, $\partial\varphi/\partial v$, and $\partial\varphi/\partial\theta$), but as for any other optimisation techniques for multivariable functions (e.g., random search, univariate search, simplex search, conjugate search directions, steepest descent, conjugate gradient methods, and Newton's methods), a 'good' starting point is required, and should be enough to get the job done. This starting point is the integer values of $[u, v, \theta]$ obtained from the preceding pixel-accuracy tracking.

Figure 7 illustrates, on an actual case, the unidimensional form of the correlation function φ along the line passing through the centres of the jumping grain in images 1 and 2 φ is plotted as a function of both the horizontal displacement u and the vertical one v as they are related to the inclination of the centre-to-centre line. The three curves correspond to linear, cubic, and quintic interpolations, and the inset is a zoom focused on the peak (where is the best probability of movement). It is quite clear that the position found depends on the interpolation function chosen; however, it is essentially the linear case that is different from higher order functions. Quintic interpolation does not seem to differ significantly from cubic one. The plots on the right part of Figure 7 also instructs us that linear interpretation

FIGURE 7 (left) Zero mean normalised cross correlation φ plotted as a function of a displacement (u and v) along the line from the initial position of a grain to its final position, for different interpolation functions (linear, cubic, and quintic). (right) Piece of interpolated greyscale maps with (bi)linear and (bi)cubic functions
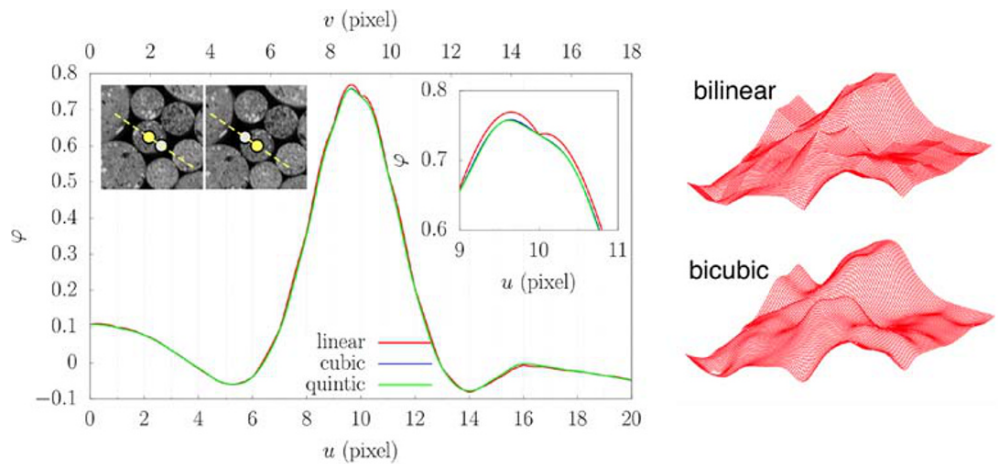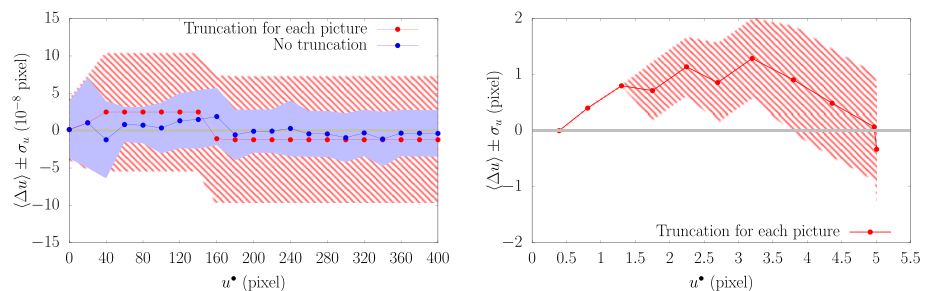

FIGURE 8 Truncation error as a function of an imposed displacement: (left) with synthetic images, (right) with photographs. The case without truncation is obtained when any image is correlated with the first image, and the case with truncation is obtained by systematically correlate pairs of successive images

induces an image mapping which is a C^0 function. This fact implies that the mapping of φ reflects this feature and suggests that linear interpolation is to be avoided.

If a deformation measurement of the tracked pattern was required rather than a rotation, the starting point could be set from an initial guess of the displacement gradients:

$$\left[u, v, \frac{\partial u_x}{\partial x}, \frac{\partial u_x}{\partial y}, \frac{\partial u_y}{\partial x}, \frac{\partial u_y}{\partial y}\right] \text{ with } \frac{\partial u_x}{\partial x} = \frac{\partial u_y}{\partial y} \simeq \cos(\theta) - 1 \text{ and } \frac{\partial u_x}{\partial y} = -\frac{\partial u_y}{\partial x} \simeq -\sin(\theta), \tag{10}$$

where $[u, v, \theta]$ are obtained from the preceding pixel-accuracy tracking.

2.4 | Truncation errors

To avoid accumulation of digital truncation errors in the sequential processing of the photographs, we took benefits of the rigid nature of the motion, that is the fact that the pattern should remain unchanged after the transformation (translation and rotation, but no—or very limited—distortion of the tracked pattern). Thus, the cross correlations are usually not performed between successive photographs ($n \to n + 1$). Instead, they are directly carried out between the first and the current photographs ($1 \to n$) so that no digital truncation error occurs. This solution can be completed efficiently if the transformation $1 \to n - 1$ is already known, while it provides an initial guess of the solution (the search zone is around the last known position, not around the reference position).

Figure 8 (left) shows the evolution of the mean error and standard deviation $\langle \Delta u \rangle \pm \sigma_u$ as the imposed displacement $u^{\bullet}$ on synthetic speckled images is increased. Despite the absence of shooting artefacts in synthetic images, errors with severest truncations (i.e., when the reference image is reset for each couple of images) are systematically higher in absolute values than the tracking without truncation (i.e., the reference image is unique and remains the same). Anyway, errors remain insignificant (less than 10^{-7} pixels) with synthetic images. With actual photographs, things are quite different as can be seen in Figure 8(right). The meaning of an error with real images is different because, unlike synthetic

images, the value of the imposed displacement is not mastered. The error is therefore in this case an increment of displacement with respect to the displacement obtained without truncation. We can see that the orders of magnitude are much more significant this time because errors of 1 or even 2 pixels can occur after processing only five pairs of photographs.

Annoyingly, if the images are of poor or nonuniform quality, the solution without truncation is not possible because the patterns would be really too altered as the sample strain evolves. It is therefore necessary to change the reference periodically, which introduces a significant accumulation of errors in the assessed displacements.

3 | PRE-PROCESSING FOR $1\gamma2\varepsilon$ TESTS

3.1 | Initial particle positions

An application has been developed specifically for the tracking of circular grains in $1\gamma2\varepsilon$ photographs. We need to start with the first photograph of the series—a 16-bits image with a single colour channel—that will serve as the reference picture (that is with reference positions of the particles) in the tracking procedure. The 2D particles, and more precisely the visible section of the Schneebeli rods, are supposed to be circular. The procedure follows a number of stages that are outlined hereafter:

- Figure 9b. Removing the parts of the image that are not of interest. This is done by setting a chosen grey-level—generally 0—to the pixels that stand outside the sample region. For the device $1\gamma2\varepsilon$, the corners of the sample frame are hinges, some of which are equipped with force sensors, which correspond to discs on the images.
- Figure 9c. Smoothing many times the sample zone of the photograph to blur the imperfections that could compromise the subsequent operations. This step is a preprocessing for the next step of segmentation. It limits the number of isolated black spots inside the grains.
- Figure 9d. Performing a segmentation to separate the particles (white or maximum grey level) from the 'voids' (black or zero grey level). The threshold used is defined from the histogram of the sample zone that have eventually been smoothed or from a greyscale profile that passes through grains and voids to be able to decide on the correct separation value.
- Figure 9e. Computing the distance map that encode, in the grey level of the pixels, the minimum distance of each pixel with a void (black) pixels. This step is the most time consuming, but it can be optimised and parallelised to finally be done in less than a minute on a modern computer.
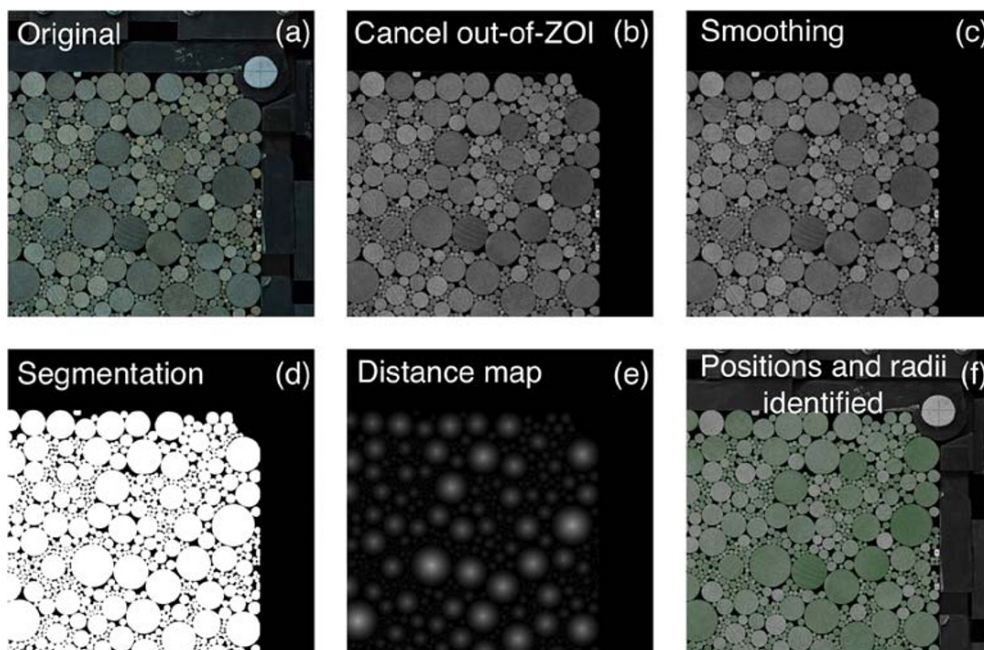

FIGURE 9   (a–f) Illustration of the different steps for the initial identification of the position and radii of the grains

- Figure 9f. Extracting the particle positions and radii in pixel units. In a nutshell, the procedure consists in iteratively finding the largest distance, which necessarily corresponds to the position of the largest radius, and then masking the grains with black (zero grey value) so that it can no longer be found. An image with superimposed green disks, similar to Figure 9f, is generated to check that the position and radius of the grains have been found in a way that is satisfactory to the eye.

Before implementing this procedure in a numerical tool coded in C++, it was necessary to spend about a day to identify the centres of 2,000 grains. The use of different applications was necessary, and multiple operator interventions were required to to correct errors 'by hand'. The current procedure identifies centres of the same number of grains in less than a minute, with substantially fewer errors.

3.2 | Accuracy of the positions of particle centres

Another source of inaccuracy when tracking a particle may come from the approximate knowledge of the position of its centre. If this position is known with a pixel resolution (integer value), the assumption of rigid motion will make the rotation of the particle not affected. On the contrary, the accuracy of displacement will be lower. Consider for example an error e on the center position, the accuracy of the tracking being ϑ (sum of maximum bias and maximum random error) for rotations. Actual inaccuracy of the displacement of the particle centre would be supplemented with $\pm 2\vartheta e$.

To limit this additional error, e should be reduced as much as possible. The determination of the particle centre r is a matter of image processing. In the current procedure, a segmentation of the image is proceeded followed by several erosions (watershed) until a single pixel remains for each particle. This last pixel is considered as the particle centre. Another strategy would be to account for the grey-level gradients at the particle borders to identify more precisely r (with subpixel accuracy). However, whatever the strategy used, the position found cannot be identified with certainty as the authentic centre. For this reason, it is arguably better to be aware of the additional error.

In TRACKER, the second strategy is used to refine the values of position and radius of each grain. The integer position of a grain—found in the preceding stage—serves as a starting point for outward-directed rays. For each half ray, the greyscale profile enables the identification, by means of a threshold value, of the distance between the temporary centre and the boundary of the grain. A number of points on the periphery of the grain then have an assigned position allowing the position and radius of a circle to be found (it best passes through these points). Because the reference position is an integer, the rest of its value is stored in an initial subpixel eccentricity e (which is a fake displacement lower than 0.5 pixel). It is actually possible to use noninteger reference positions, but doing so requires to perform a grey interpolation on the reference image in addition to the interpolation of the target image. Although doable, we did not choose this solution because it required significant changes in the source code, and it could have led to a slowing down of the correlation algorithm which is sometimes also employed to track movements on continuous materials. The integration of this rather subtle correction was thus postponed to the post-processing phase. To benefit from the improved accuracy associated with the initial eccentricity, it is necessary to take into account the initial displacement in postprocessing by defining any position as $r + u + R(\theta) \cdot e$; see Figure 10.
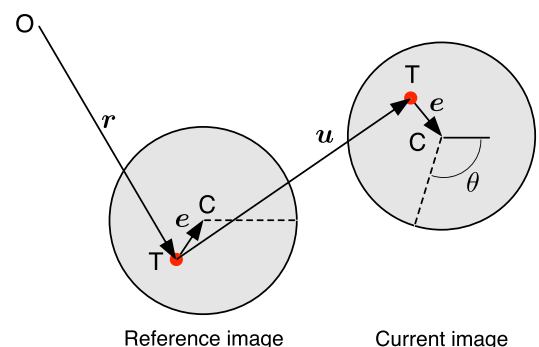

FIGURE 10 Illustration of the consideration of subpixel eccentricity e (oversized in the sketch)

Reference image   Current image

4 | CORRELATION ERRORS WITH IDEALISED IMAGE

In a previous study, TRACKER accuracy was estimated from synthetic images provided by[12]. These synthetic images simulated the rigid motion of a speckle texture captured by an idealised camera (perfect sensors with fill-factor equals to 1). In the present study, two series of images were independently generated with different kinds of transformation: pure horizontal translation and homogeneous vertical straining due to the increasing displacement of the points on either side of the vertical centre line. The patterns for these accuracy tests was generated by Perlin noise, and the achievement of best qualities was not sought for these synthetic images. The objective here is to test the dependence of the subpixel accuracy with respect to the degree of the polynomial function (bilinear or bicubic) used to interpolate the image and the size (pixels) of the squared patch.

4.1 | Pure translation

An imposed displacement $u^\bullet$ was varied from 0 to 1 pixel with a step of 0.02 pixel. We compared $u^\bullet$ with the corresponding displacement u_i of each of the 400 points i followed to determine their displacement error: $\Delta u_i = u^\bullet - u_i$. The arithmetic mean of Δu_i defined the systematic error (or bias):

$$\langle \Delta u \rangle = \frac{1}{n} \sum_{i=1}^{n} \Delta u_i. \tag{11}$$

In addition, the standard deviation (random error) σ_u was computed:

$$\sigma_u = \sqrt{\frac{1}{n-1} \sum_{i=1}^{n} (\Delta u_i - \langle \Delta u \rangle)^2}. \tag{12}$$

As for other accuracy studies[13,14], displacements were evaluated at all positions of a regular square grid in the initial image, with a pitch such that correlation windows at adjacent positions do not overlap, ensuring the statistical independence of the corresponding errors.

Figure 11 shows the bias and the random error obtained for imposed translations $u^\bullet$ in the horizontal direction. In this representation, standard deviations are indicated with error bars, but all data are also shown (which is never done to our knowledge). This allows a more correct view of the dispersion of random errors. The observation that may seem most disturbing is that the bias is less close to zero when the order of the interpolation function is increased. This is in fact the price to pay so that overall random errors are lowered. It can be seen, however, that biquintic interpolation does not bring any improvement, either for bias or for random errors. For a given size of patterned patch, random errors are smaller for bicubic than for bilinear interpolations.

Focusing now on the size of the tracked patch, it can be observed that although the bias is almost independent of the pattern size, the random error is decreased with the size of the pattern (whatever the degree of the interpolation function). These results are consistent with those of Bornert et al.[14] or Dupré et al.[13]. Here, for this particular synthesised pattern, the worst error in translation for the largest pattern (33×33) is in the order of 2×10^{-2} pixel regardless of the interpolation function, even if the error seems to be improved from the point of view of the standard deviation. For smaller patch sizes, bicubic interpolation has a more definite gain over overall accuracy. Finally, the bilinear interpolation seems a sound choice for large patterns (in terms of accuracy in the displacement, with the use of synthetic images), while cubic interpolation should be more appropriate when the size of the patterns is limited (e.g., by grain sizes).

This analysis, even if it is carried out on synthetic images without aberration, makes it possible to realise the equal relevance of the quality of the pattern and of the size of the tracked patches. The goodness of a pattern must therefore be judged also in terms of the usable patch size. However, it is helpful to consider that a given pattern, interpolated with any function, will systematically give better results with larger patterns as suggested by Figure 12.

4.2 | Homogenous straining

When soft-grain movements are tracked from the image of $1\gamma2\varepsilon$ loaded granular sample or some textured points are tracked within a close view of particles in contact, noticeable finite strain can occur. It is then pertinent to wonder to
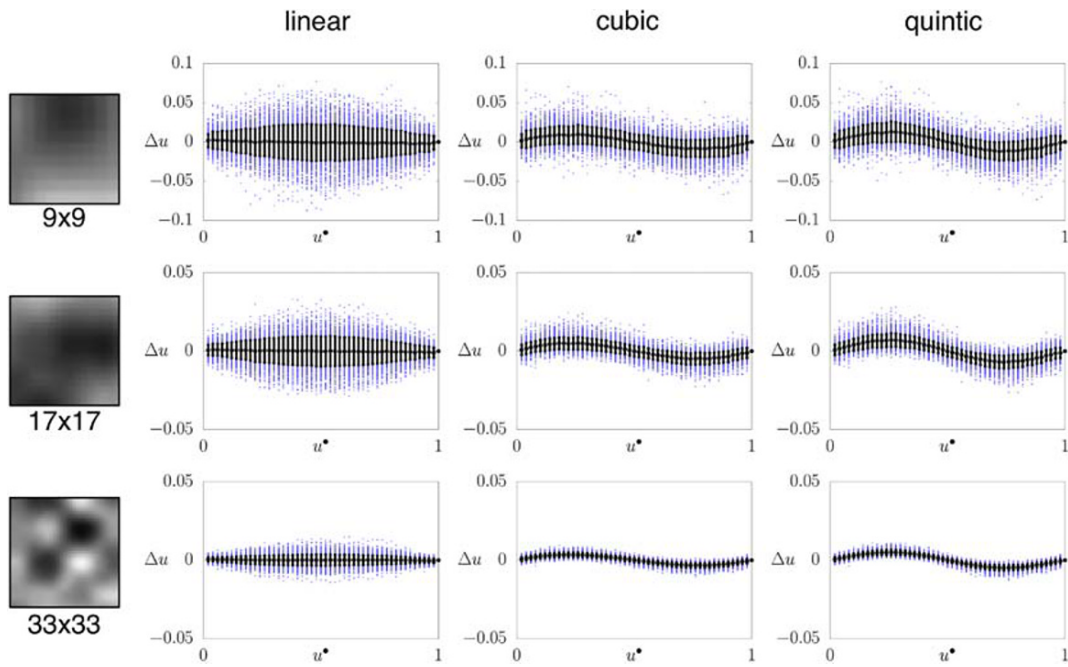
FIGURE 11 Displacement errors Δu as a function of imposed displacement $u^{\bullet}$ for three sizes of tracked patches and increasing degree of the interpolation function. The blue dots are the raw displacement errors, the black dots are the bias errors (arithmetic means), and the error bars are the random errors (standard deviations)

FIGURE 12 Maximum random error as a function of the inverse of pattern size (N_m is the number of pixels in a patch)
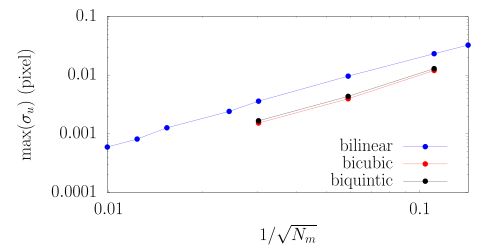
what extent such deformations can alter the measured displacements. To address this question, synthetic images have been generated, still with Perlin noise, by imposing a vertical swelling strain $\varepsilon^{\bullet}$ ranging from 0 to 0.1 with increments of 2×10^{-3}. By knowing the position of vertical middle r_x^0 of the images, the expected vertical displacement of for any coordinate r_x is $u(r_x) = \varepsilon^{\bullet}(r_x - r_x^0)$. This allowed us to define the error as

$$\Delta u(r_x) = u(r_x) - \varepsilon^{\bullet}(r_x - r_x^0) \tag{13}$$

Figure 13 shows this error Δu as a function of the imposed strain $\varepsilon^{\bullet}$ for three sizes of pattern and when using bilinear or bicubic interpolations. Three observations can be made by looking at these plots: (a) the errors are not affected by the choice of interpolation function, (b) random errors are linearly increased with increasing deformation, and (c) increasing the size of the pattern magnifies the random errors in case of finite strain, but the bias does not seem affected. This last remark might seem counterintuitive, but is in fact expected because a given elongation generates more change over a large area than over a small area.

It must be clarified that it would be sufficient to track the deformations of the correlation patch to annihilate the errors—by optimising the parameters of Equation (6) for instance. However, in the applications we are interested in, the deformations within the tracked speckles on grains are not at all homogeneous because they are localised at the vicinities of contact zones. It is then convenient not to include the edge of the grains in the thumbnail which is tracked.
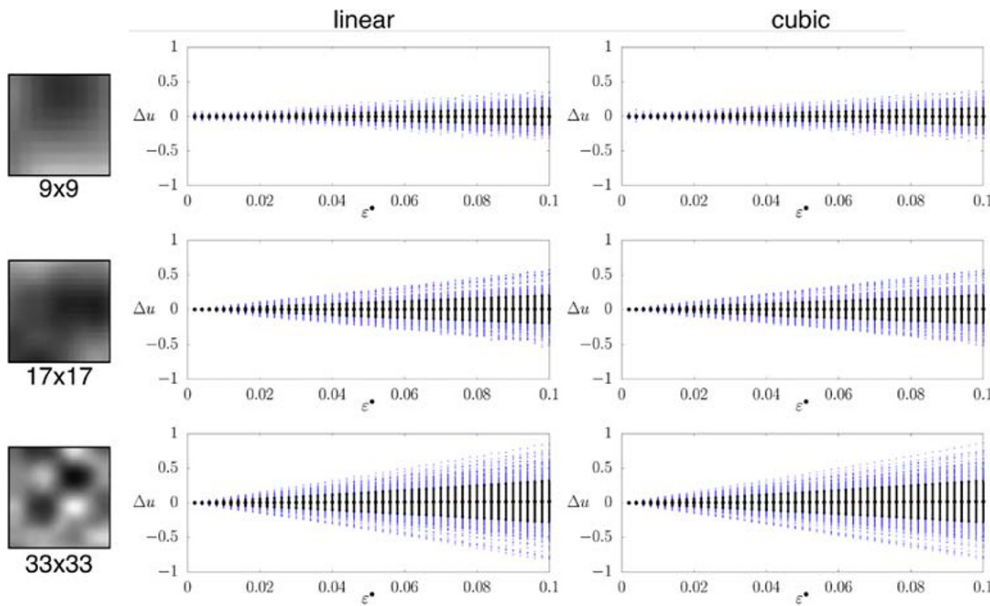
FIGURE 13 Displacement errors Δu as a function of imposed vertical strain $\varepsilon^\bullet$ for three sizes of tracked patches and increasing degree of the interpolation function. The blue dots are the raw displacement errors, the black dots are the bias errors (arithmetic means), and the error bars are the random errors (standard deviations)

5 | CORRELATION ERRORS WITH REAL IMAGES

For real images of moving bodies coming from a digital camera, the estimate of the error for displacements and rotations is rather complex because it implies a perfect mastery of the whole measurement chain. In particular, the digital images are captured on a Bayer matrix containing ordered red, green, and blue cells (two greens for one red and one blue) so that each pixel of the produced image will have three colour channels. Each channel is built from a mixing recipe, called debayering or demosaicing, supposed to produce the best image according to a criterion of its choice. With DIC, nobody cares about the subjective beauty of the image, but the accuracy of the tracked displacements is the only feature required. Different tests (not shown here) led us to conclude that the simplest debayering formula (filter) did the trick and no clear trend could be drawn from the choice of filter. So simplest linear filter is systematic employed. The best solution would be to choose a camera without bayer matrix (raw images directly in greyscale).

The accuracy study for real images is here performed over photographs of an assembly of about 2,000 rods submitted to a quasi-static shearing like the one shown in Figure 1($\dot{\gamma} = 8.2 \times 10^{-5}$ s^{-1}). The camera was a Nikon D3X (24.4 MPixels) associated with a lens of fixed focal length of 180 mm(Tamron SP AF 180 mmF/3.5 Di LD[IF] MACRO 1:1). We used also, in other tests, a Nikon D850 (42.5 MPixels) associated with a lens of fixed focal length of 180 mm, or a PhaseOne IQ180 digital back (80 MPixels) associated with another fixed focal lens (Schneider Kreuznach, AF 240 mmLS f/4.5 IF). Both lens correspond to long focal lengths, and they were chosen to limit the field of view because the shortest values make the rod axis appear more parallel on the photographs. The first lens induces a quite homogeneous distortion while the second concentrates the distortions at the photograph edges. The camera was always placed relatively far from the shot sample. Along the shear test, pictures were shot every 5 s. The correction of optical distortion which requires a specific procedure (see Appendix D1) was either performed or not to examine the consequences. Figure 14 shows the bias and random error of the measured variations of length ΔS, placed within the largest grains, throughout the shear test. It is shown that, in a first phase up to $\gamma \simeq 0.08$, $\Delta S \leq 0.1$ pixel, with an average of $\langle \Delta S \rangle \leq 0.01$ pixel. The benefits of the distortion correction is clear because the mean value $\langle \Delta S \rangle$ remains close to zero, whereas a very clear drift is observed without correction. In a second phase of the shear for γ above 0.08, the error increased for an unnoticed reason, and the drift was worse. We can see that the distortion correction does not really improve the random errors on ΔS, but the bias is on the other hand very significantly reduced.

To perform a PIT, any other lens could be used, and if the focal length is particularly short, the need of correction of the distortion is obvious. But correction is always necessary if a high degree of measurement accuracy is required. The plots of Figure 14 shows that an average error of about 0.02 pixels can be avoided by correcting the distortion, which is nevertheless only slightly perceptible to the eye (after meticulous inspection, it is possible to notice that the circularity of the grains is only slightly affected at the edge of the image). Of course, the measurement of the error ΔS is based on the assumption that the grains do not deform noticeably. We carried out compression tests on grain pairs with forces

FIGURE 14 Mean and standard deviation of length variations in the course of shearing: (red) without and (blue) with correction of distortion
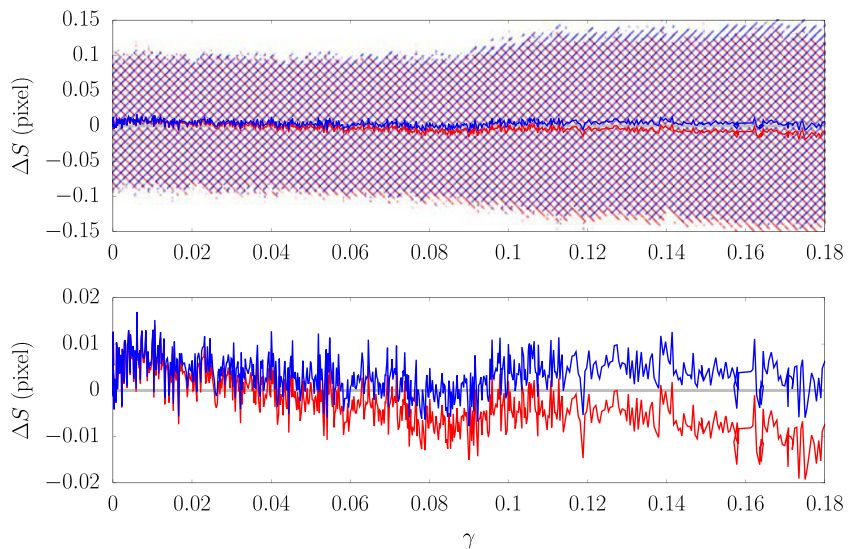
40 times greater than the typical forces in the experiments. With about 2,000 pixels per diameter and a well-established correlation tool[15], we found that both spherical strains and deviatoric strains remained extremely small inside the grains, and concentrated only at the vicinity of the contact.

The objective of distortion correction as explained in Appendix D1 is to guarantee that the same amount of displacement is quantified in the least fluctuating way, whatever the location in the image. In other words, the optimisation of the equiprojectivity criterion on a rigid body movement covering the entire image seems to us self-sufficient.

6 | EXAMPLE OF USING THE PIT

The use of the PIT technique is presented here through one example that requires a high accuracy of the tracked displacements: the measurement of displacement fluctuations in a sheared granular assembly[8,16,17]. The work of Misra and Jiang [18] is, to our best knowledge, the first scientific reference showing a map of experimentally measured fluctuations. In this paper, the authors show some maps of deviations from linear displacement fields (we usually call it 'displacement fluctuations'). It has been shown, in recent studies, that a fluctuation analysis requires consideration of a deformation window in the definition of a displacement fluctuation. Typical observations of 'turbulent-like' vortices in the fluctuation field—called 'clustering patterns' in[18]—are sensitive to the macroscopic deformation increment under consideration. In Misra and Jiang[18], the increment is not specified, but it should be relatively large and the statistical distribution of the components of displacement fluctuation should certainly be close to a bell-shaped Gaussian distribution. For decreasing deformation windows, the Gaussian distribution of normalised fluctuations tends to broaden and correspond more closely to a q-Gaussian distribution, while nonnormalised fluctuations become smaller and smaller and therefore difficult to measure on digital images. The need for precision was particularly felt in the analysis of these small fluctuations. Other applications that require a good accuracy are the detection of contacts and the assessment of elastic contact forces as in the work of Tolomeo et al.[19]. The purpose here is to show a practical use of the technique for measuring displacement fluctuations rather than enter into too deep understanding of the underlying features as originally described in Radjai and Roux[20].

Displacement fluctuations in granular materials are a direct manifestation of grain rearrangement, which is the basic mechanism of irreversible deformation; fluctuations play a similar role as dislocations in crystals. The analysis of fluctuations is a hot topic in granular physics and mechanics, which is almost entirely investigated by means of numerical discrete element model (DEM). To assess the fluctuations of displacement in the course of deformation, two possible displacements of each grain during a shear strain increment $\Delta\gamma$ are considered. The first is the actual displacement vector $\delta r(\gamma, \Delta\gamma)$, which depends on both the size of the strain increment $\Delta\gamma$ and the level of shear strain γ at the beginning of the increment. The second displacement vector $\delta r^\star(\gamma, \Delta\gamma)$ is the displacement dictated by a homogeneous (*affine*) continuum strain field, that is, the displacement that the grain's centre would have if it
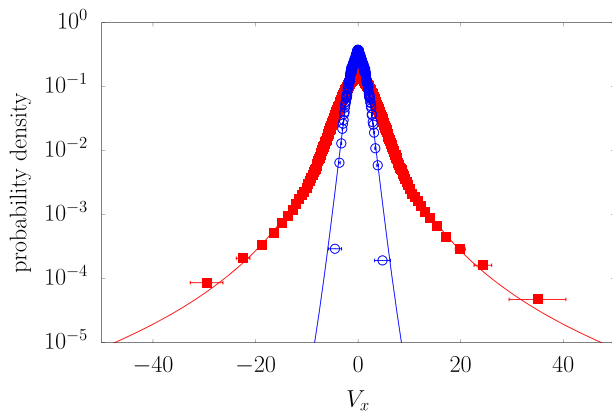
FIGURE 15 Probability density function of the horizontal component of the displacement fluctuation V_x for two values of shear strain increment $\Delta\gamma$ (red squares) in case of a simple shear test on a sample of around 2,000 rods in the device $1\gamma2\varepsilon$. Symbols are PIT inputs (red squares: 2×10^{-3}, blue circles: 10^{-1}). The curves are q-Gaussian fits (red: $q = 1.51$, blue: $q = 1.16$)

moved as a material point within a continuum. The fluctuation of the displacement is defined as the difference between these two displacements:

$$\boldsymbol{u}(\gamma, \Delta\gamma) = \delta\boldsymbol{r}(\gamma, \Delta\gamma) - \delta\boldsymbol{r}^\star(\gamma, \Delta\gamma). \tag{14}$$

Displacement fluctuations can be conveniently normalised by dividing $\boldsymbol{u}(\gamma, \Delta\gamma)$ by the product $\Delta\gamma\langle d \rangle$ (where $\langle d \rangle$ is the mean diameter of the grains), which can be interpreted as the average displacement of the grains in the strain increment $\Delta\gamma$. This normalised fluctuation, $\boldsymbol{V}(\gamma, \Delta\gamma) = \boldsymbol{u}(\gamma, \Delta\gamma)/(\langle d \rangle\Delta\gamma)$, can also be interpreted as a local strain fluctuation which is in turn divided by the size of the global strain increment $\Delta\gamma$.

Figure 15 shows the probability density function (pdf) of the horizontal component V_x of normalised fluctuations for two different shear strain increments $\Delta\gamma$. Whatever the value of $\Delta\gamma$, the space average of this measure of fluctuations is zero, which is to be expected for a homogeneous deformation. A key observation is that the pdf curve exhibits a wider range of fluctuations with decreasing $\Delta\gamma$. An extremely accurate measurement of fluctuations is necessary to be able to define their statistical distribution. Besides, in this case, it is not a simple Gaussian but rather a q-Gaussian distribution function as introduced by Tsallis[21]. The evolution of such parameters (here the parameter q of the q-Gaussians) gives valuable clues on fundamental aspects of the plasticity of granular media, and it is therefore absolutely essential to push the technique to its limits, in terms of accuracy, in order to obtain the most useful measures possible.

7 | CONCLUSIONS

A new parallelised code called TRACKER dedicated to the tracking of nonsmooth trajectory of particles was presented. The software uses the DIC technique coupled with geometrical rules to dramatically increase the efficiency of the cross correlations. These tremendous improvements allow for an accurate and fast assessment of the grain kinematics without 'losing' any of them. The proposed strategy, named PIT, was successfully used to exhibit very interesting features of particle fluctuations in sheared dense granular systems[8,9], but also to extract some components of fabric data within a granular packing, or assess the contact forces between the grains[19].

Many strategies have been implemented to achieve the best possible accuracy: rescue procedure, interpolation function, correction of distortion and of small parallax flaws, reduction of truncation errors, reduction of errors due to the eccentricity of the tracked points, and correction of parasitic movements (not shown in this paper). Tests were made on computer-generated images in pure displacement and also in homogeneous deformation. Then tests with real images were also carried out. The finding is that it is difficult to do better in terms of the tracking algorithm. The other possibilities of improvements do not relate directly to the technique but to the quality of the patterns and digital images. So, to further improve accuracy, work must be done on image acquisition and pattern quality.

ORCID

Vincent Richefeu 🄳 https://orcid.org/0000-0002-8897-5499

REFERENCES

H. Joer, J. Lanier, J. Desrues, E. Flavigny, *Geotechnical Testing Journal* **1992**, *15*(2), 129.

G. Schneebeli, *Comptes Rendus de l'Academie des Sciences* **1956**, *243*(1), 125.

F. Calvetti, G. Combe, J. Lanier, *Mechanics of Cohesive Frictional Materials* **1997**, *2*(2), 121.

E.-M. Charalampidou, G. Combe, G. Viggiani, J. Lanier, *in Powders and grains*, (Eds: S. L. Masami Nakagawa), AIP Publishing, Golden, Colorado **2009**.

L. Sibille, F. Froiio, *Granular Matter* **2007**, *9*, 183.

T. C. Chu, M. A. Sutton, W. F. Ranson, W. H. Peters, *Experimental Mechanics* **1985**, *25*(3), 232.

S. A. Hall, D. Muir Wood, E. Ibraim, G. Viggiani, *Granular Matter* **2010**, *12*(1), 1.

V. Richefeu, G. Combe, G. Viggiani, *Géotechnique Letters* **2012**, *2*, 113.

G. Combe, V. Richefeu, G. Viggiani, S. A. Hall, A. Tengattini, A. P. F. Atman, Tracker: A particle image tracking (PIT) technique dedicated to nonsmooth motions involved in granular packings, in Powders and grains **2013**.

M. Powell, *The Computer Journal* **1964**, *7*, 155.

W. H. Press, S. A. Teukolsky, W. T. Vetterling, B. P. Flannery, *Numerical recipes 3rd edition: The art of scientific computing*, Cambridge University Press **2007**.

J.-J. Orteu, D. Garcia, L. Robert, F. Bugarin, A speckle-texture image generator, in Proc. SPIE 6341, Speckle06: Speckles, From Grains to Flowers, 63410H **2006**.

J. C. Dupré, M. Bornert, L. Robert, B. Wattrisse, Digital image correlation: Displacement accuracy estimation, in Proceedings of the 14$^t h$ international conference on experimental mechanics **2010**, 31006.

N. Bornert, F. Brémand, P. Doumalin, J.-C. Dupré, M. Fazzini, M. Grédiac, F. Hild, S. Mistou, J. Molimard, J.-J. Orteu, L. Robert, Y. Surrel, P. Vacher, B. Wattrisse, *Experimental Mechanics* **2009**, *49*(3), 353.

P. Vacher, S. Dumoulin, F. Morestin, S. Mguil-Touchal, *Journal of Mechanical Engineering Science* **1999**, 811.

G. Combe, V. Richefeu, M. Stasiak, A. Atman, *Physical Review Letter* **2015**, *115*, 238301.

L. Viallon-Galinier, G. Combe, V. Richefeu, A. P. F. Atman, *Entropy* **2018**, *20*(11), 862.

A. Misra, H. Jiang, *Computers and Geotechnics* **1997**, *20*, 267.

M. Tolomeo, V. Richefeu, G. Combe, J.-N. Roux, G. Viggiani, *International Journal of Solids and Structures* **2019**, *187*, 48.

F. Radjai, S. Roux, *Physical Review Letter* **2002**, *89*, 064302.

C. Tsallis, *Journal of Statistical Physics* **1988**, *52*, 479.

M. A. Sutton, J.-J. Orteu, H. Schreier, *Image correlation for shape, motion and deformation measurements: basic concepts, theory and applications*, Springer **2010**.

Y. Su, Q. Zhang, X. Xu, Z. Gao, *Optics and Lasers in Engineering* **2016**, *86*, 132.

D. C. Brown, *Photogrammetric Engineering* **1966**, *32*(3), 444.

APPENDIX A: QUALITY OF THE PATTERN SIGNALS

To track particles on a 2D image, the speckles that are photographed on the visible faces of the rods will not be of equal quality with respect to the accuracy of the displacement measurements. This accuracy will of course depend on the speckle definitions but also on the pattern sizes. The presence of well marked greyscale gradients also seems to play an important role in the reliability of particle tracking. The literature essentially proposes two types of indicators to quantify the quality of a pattern: one based on its morphology while the other is based on its similarity to itself.

In the use of the tool TRACKER for our research applications, we have chosen two measures of the pattern quality based on its self-similarity: its Fourier transform and the mean Pearson correlation of any small patches within the

pattern as a function of their distance (self-correlations). Figure A1 shows, for different pattern patches of size 101×101 (extracted from a dataset of the DIC challenge, https://sem.org/dicchallenge), the corresponding Fourier transforms and also a plot of self-correlations. It is important to notice that the plot of self-correlations is not really identical to what is described in the literature, that are the auto-correlograms (see, e.g.,[14,22]); but the underlying idea is exactly the same. The Fourier transform enables the identification of a characteristic size of a spot (frequency bandwidth), and the self-correlation plot makes it possible to define a distance δ_{self} reached for a half-reduced correlation.

What must be noted is that none of these indicators is capable of discriminating with absolute certainty a good from a bad pattern. Nevertheless, it provides access to quantifiable parameters which, with a certain amount of expertise, can be used to guide the production of painted speckles. Also, when a tracking does not go well, the explanation may come from a pattern which is too poor in quality, and the assessment of a quality indicator is quite useful in this case. To prevent relying too much on the expertise of operators alone, the use of several indicators is advisable.

Table A1 shows the values of four indicators: (a) the radius δ_{fft} (inverse of the spatial frequency) of the Fourier transform spot, (b) the distance δ_{self} of half-reduced correlation in the self-correlation diagram (obtained with a 5-by-5 patch), (c) the distance δ_{auto} of half-reduced correlation in the classical auto-correlogram, and (d) the mean intensity gradient δ_{mig} as defined in Su et al.[23]. This latter indicator provides information on the greyscale gradients that would be extreme in the case of a pattern made of black dots every 2 pixels on a white background (upper limit), but also in the case of a mono-level grey image (lower limit). It is defined as follows:

$$\delta_{\text{mig}} = \frac{1}{WH} \sum_{i=1}^{W} \sum_{j=1}^{H} \sqrt{(\nabla_x I_{ij})^2 + (\nabla_y I_{ij})^2}, \tag{A1}$$
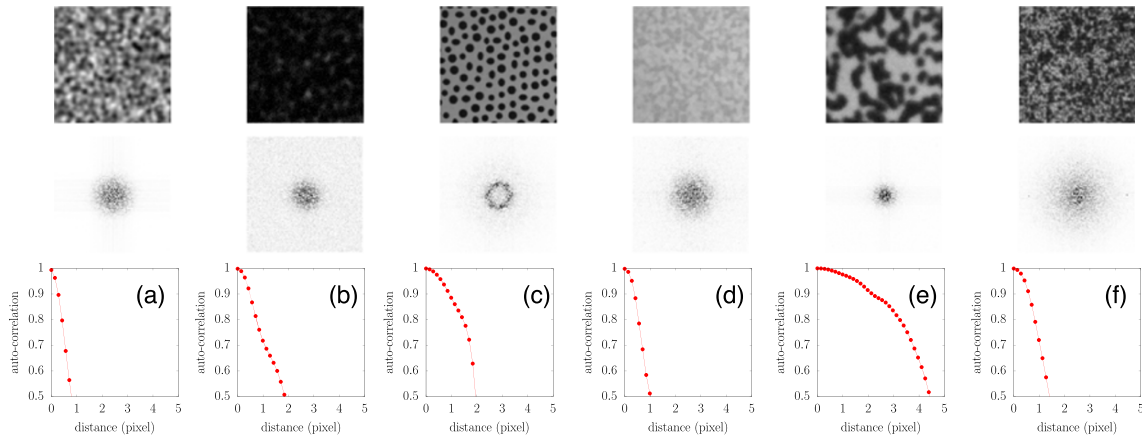

FIGURE A1 (a–f) Quality indicators of some speckle patterns from the DIC Challenge Dataset (https://sem.org/dicchallenge). Top line: 100 × 100 digital image with different patterns. Middle line: the corresponding Fourier transform (power increases from white to black). Bottom line: mean Pearson correlation of 101-by-101 patches on the same image as a function of their distance. This latter plot of self-correlations, somehow analogous to a auto-correlograms, shows the distance reached for a half-reduced correlation; this distance, noted δ_{self} and expressed in pixel, defines a measure of the pattern quality (more accurate DIC may be obtained with smaller value of δ_{self})

TABLE A1 Four indicators of pattern quality, δ_{fft}, δ_{self}, δ_{auto}, and δ_{mig}, computed on 101-by-101 patches. The smallest values are meant to provide the best quality, except for δ_{mig} which should not be too small nor too high

Pattern	δ_{fft}	δ_{self}	δ_{auto}	δ_{mig}
(a)	0.07	0.8	3.6	0.12
(b)	0.08	1.8	2.1	0.06
(c)	0.08	1.9	2.1	0.17
(d)	0.05	1.0	2.9	0.04
(e)	0.10	4.4	4.7	0.11
(f)	0.08	1.4	1.9	0.16

where ∇_x and ∇_y stand respectively for the grey-level gradient in the x- or y-direction, W and H are the width and height of the cropped zone in the pattern image.

From Figure A1 and Table A1, some interesting observations can be drawn. (b) and (c) appear to be of comparable quality with the exception of the criterion on δ_{mig} that exhibits extreme values. This is due to an image that is too dark for (b) (which flattens the gradients) and too high in contrast for (c) (characterised by very high gradients at the transitions between white and black). For this reason, pattern (a) or (f) would be a better choice to obtain greater measurement accuracy because the indicators are on the whole more beneficial. It is instructive to see that patterns (b) and (d) are not as bad as they seem and that on the contrary, pattern (e) is clearly a wrong choice even if everything suggests otherwise. It is nevertheless unfortunate to realise that some empiricism is still needed in the definition or choice of sound patterns. It is conceivable that a useful quantification of the quality of a pattern, with respect to the accuracy it would allow, can be established by combining indicators δ_{auto} and δ_{mig} as a unique marker. For instance, $\delta_{\mathrm{auto}}^{\alpha}\delta_{\mathrm{mig}}^{\beta}$, where α and β have to be searched by trial and error, would be a candidate. However, this possibility has not yet been studied.

APPENDIX B: PARALLELISATION

Most numerical procedures during the tracking, including rescue procedures that are carried out subsequently, can be applied for each grain regardless of the others. This makes the parallelisation very easy to implement in a context of shared memory (OpenMP). Figure B1 shows the relative gain obtained on two computers as a function of the number of threads requested. The relative gain G is defined as the duration (CPU time) obtained with a unique thread divided by the duration obtained with N_{th} threads. The first computer is a MacBook Air laptop with a single dual-core Intel Core i7, and the second one is a pair of Intel Xeon processor E5-2683-V4 with 40 MB of cache memory (hyperthreading no activated); the number of available threads is 2 and 32, respectively. Regardless of the hardware used, a linear increase in relative gain is observed up to the number of available threads. Above this point, the gain may drop very slightly, but no longer varies.

Still in the same plots, the relative gain G was calculated separately for the part of the algorithm dedicated to the pixel-by-pixel search and for the subpixel optimisation. The scalability of these two parts is similar, but two features are particularly intriguing in the plot with 32 threads: (a) the rate $\Delta G/\Delta N_{\mathrm{th}}$ of relative gain is lowered from four threads upwards, and (b) the odd number of N_{th} induces lower gains, perhaps because of memory management.

APPENDIX C: DISTORTION AND PARALLAX OF ACQUIRED IMAGES

The projection of the physical objects onto the photographs is not strictly rectilinear because of (*i*) imperfect perpendicularity of the optical axis with the plane of the assembly (parallax defect) and/or (*ii*) the lens characteristics (distortion).

The deviation of this projection, called *distortion*, makes the relative distances not uniform over an image. The displacement measurements are thus affected by the distortion as a function of the position of the tracked pixel, and it is
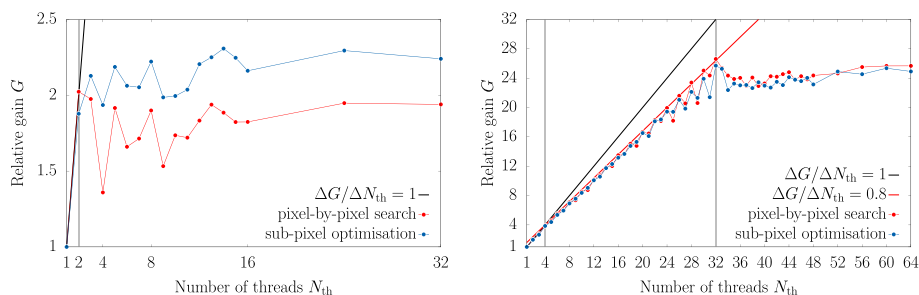

FIGURE B1 Gain obtained with respect to the calculation with just one thread as a function of the number of threads used (this number can be larger than the available cores): (left) for laptop with single dual-core processor, (right) for a computer with four octo-core processors. The red connected dots are for the time required for the pixel-by-pixel search phases, and the blue connected dots refer to the times required for the sub pixel search phases
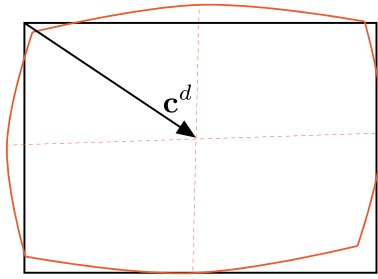
FIGURE C1   Illustration of what the frame of a distorted image may look like. The centre of distortion c^d is shown, but the other parameters are not explicitly represented

increased for large displacements (that can be of the order of third the photograph width for a shear test). To overcome this problem, a procedure based on a decentring distortion model that account for both optical distortion and parallax defect has been implemented. This model makes use of a set of parameters $\{\phi_d\}$ to determine the actual (i.e., undistorted) position $\boldsymbol{r}^\star = [r_x^\star, r_y^\star]$ of a pixel from its current position $\boldsymbol{r} = [r_x, r_y]$. The nonlinear operator $\mathcal{D}$ used is the optical distortion model of Brown[24] that reads

$$\boldsymbol{r}^\star = D_{\{\phi_D\}}(\boldsymbol{r}) \equiv \begin{pmatrix} r_x^\star & = r_x + d_x(K_1\rho^2 + K_2\rho^4 + K_3\rho^6) + (P_1(\rho^2 + 2d_x^2) + 2P_2 d_x d_y)(1 + P_3\rho^2) \\ r_y^\star & = r_y + d_y(K_1\rho^2 + K_2\rho^4 + K_3\rho^6) + (P_1(\rho^2 + 2d_y^2) + 2P_2 d_x d_y)(1 + P_3\rho^2) \end{pmatrix}, \tag{C1}$$

where $\boldsymbol{d} = [d_x, d_y]$ is the position $\boldsymbol{r}$ relative to the distortion center $\boldsymbol{c}^d = [c_x^d, c_y^d]$; $\rho = |\boldsymbol{d}|$; K_1, K_2, and K_3 are radial distortion coefficients; and P_1, P_2, and P_3 are tangential distortion coefficients. Figure C1 gives an illustration of what a distortion may look like. The set of distortion parameters $\{\phi_D\}$ is thus constituted by the eight-parameter set $\{c_x^d, c_y^d, K_1, K_2, K_3, P_1, P_2, P_3\}$.

To correct distortions and parallax simultaneously, a speckled rigid pattern that fill the photograph is shot in two different positions. A number of pixels that are dispatched on the whole image to define a set of positions $\{\boldsymbol{r}_i\}$ are then tracked to obtain their corresponding displacements $\{\boldsymbol{u}_i\}$. Corrected positions $\{\boldsymbol{r}_i^\star\}$ and displacements $\{\boldsymbol{u}_i^\star\}$ can be generated by applying the operator $\mathcal{D}_{\{\phi_D\}}$. Finally, the parameters of distortion must lead to a displacement field that best meet the equiprojectivity criterion. This parameters should thus minimise the following function:

$$\Psi_{\{\phi_D\}}(\{\boldsymbol{r}_i\}, \{\boldsymbol{u}_i\}) = \sum_i \sum_{j>i} \left| (\boldsymbol{r}_j^\star - \boldsymbol{r}_i^\star)(\boldsymbol{u}_j^\star - \boldsymbol{u}_i^\star) \right|, \tag{C2}$$

where $\{\phi_D\}$ is the set of parameters to be optimised. Usually, before starting the optimisation, all parameters are set to zero, except the distortion center which is initially placed in the center of the image. In addition, the equiprojectivity criterion is in practice verified on several image pairs p, which implies minimising a sum of errors $\sum_p \Psi_p$ given by Equation (C2).

At this stage, it can be noticed that Equation (1) actually performs a distortion operation rather than an inverse-distortion one. However, we performed a number of tests of distortion correction by using the reverse operation $\mathcal{D}^{-1}$ and found that the correction quality was of the same order; $\Psi_{\{\phi_D\}} \simeq \Psi_{\{\phi_{D-\infty}\}}$. Because it is precisely the only condition that matters, the form of Equation (C2) is generally used (as usually done in the literature).

To feel the consequences of the distortion on the displacement errors obtained, let us note a corrected displacement $\boldsymbol{u}^\star$ as the difference of corrected position of a given point on a pair of photographs $\boldsymbol{R}^\star - \boldsymbol{r}^\star$. Introducing the distortion operator, the relative displacement error expresses as

$$\boldsymbol{u}^\star \oslash \boldsymbol{u} = (D(\boldsymbol{r} + \boldsymbol{u}) - D(\boldsymbol{r})) \oslash \boldsymbol{u}, \tag{C3}$$

where the symbol $\oslash$ (Hadamard division) has to be understood as a component wise division. It is interesting to notice from this expression that

(1)   the relative displacement error depends on the displacement magnitude itself,

(2) the error on a displacement is the distortion gradient in the limit of small displacements,

(3) for a nonzero constant gradient, the greater the amplitude of the displacement, the greater the error.

With these remarks in mind and knowing that the movement of the grains is very erratic in a deforming granular assembly, it is very important to correct the distortion of the images to obtain the best possible accuracy.