An extension of digital volume correlation for multimodality image registration

Erika Tudisco, Clément Jailin, Arturo Mendoza, A Tengattini, E Andò, Stephen Hall, Gioacchino Cinno Viggiani, Francois Hild, Stéphane Roux

An extension of digital volume correlation for multimodality image registration

E Tudisco[1], C Jailin[2], A Mendoza[2], A Tengattini[3,4], E Andò[3], S Hall[5], G Viggiani[3], F Hild[2] and S Roux[2]

[1] Division of Geotechnical Engineering, Lund University, Sweden
[2] LMT (ENS Paris-Saclay/CNRS/Univ. Paris-Saclay), Cachan, France
[3] Univ. Grenoble Alpes, CNRS, Grenoble INP‡, 3SR, F-38000 Grenoble, France
[4] Institut Laue-Langevin, 71 Avenue des Martyrs, 38000 Grenoble
[5] Division of Solid Mechanics, Lund University, Sweden

E-mail: clement.jailin@ens-paris-saclay.fr

Abstract. The question of registering two images (or image volumes) acquired with different modalities, and thus exhibiting different contrast, at different positions is addressed based on an extension of global Digital Image (or Volume) Correlation. A specific comparison metric is introduced allowing the signature of the different phases to be related. A first solution consists of Gaussian mixture to describe the joint distribution of gray levels, which not only provides a matching of both images, but also offers a natural segmentation indicator. A second 'self-adapting' solution does not include any postulated *a priori* model for the joint histogram and leads to a registration of the images based on their initial histograms. The algorithm is implemented with a pyramidal multiscale framework for robustness. The proposed multiscale technique is tested on two 3D images obtained from X-ray and neutron tomography respectively. The proposed approach brings the two images to coincidence with a sub-pixel accuracy and allows for a "natural" segmentation of the different phases.

Keywords: Digital image correlation; Image fusion; Image registration; Neutron tomography; X-ray tomography

Submitted to: *Meas. Sci. Technol.*

1. Introduction

Different imaging modalities can be be used to gain information on material properties and structures, but the information gained from each modality will be different due to the different sensitivities of the measurement approach used. This leads to different contrasts in images of the same object, but also a greater richness of information. To benefit from multiple modal imaging, it is essential to obtain a local characterization

‡ Institute of Engineering, Univ. Grenoble Alpes

from each image for the *same* material point. However, different modality images are generally not acquired (or reconstructed) in the same frame. Hence, rigid body motions have to be considered to make the two frames coincident. Additionally, the resolutions of the images may be different, which means that the transformation to map the images onto one another should be enriched with a possible dilation/contraction. The purpose of the present study is to identify the transformation $\boldsymbol{F}$ (*i.e.,* rigid body motions and isotropic dilation, corresponding to 4 or 7 parameters in 2D or 3D, respectively) that enables the optimal registration of two images acquired with different modalities.

As an illustration, Figure 1 shows two horizontal slices extracted from 3D volume images of the same cylindrical specimen of a sandstone imaged by X-ray and neutron tomography. The contrast is different between the two images, although some specific (and matching) patterns can be identified. This 3D example will be used to illustrate and validate the proposed image registration procedure.

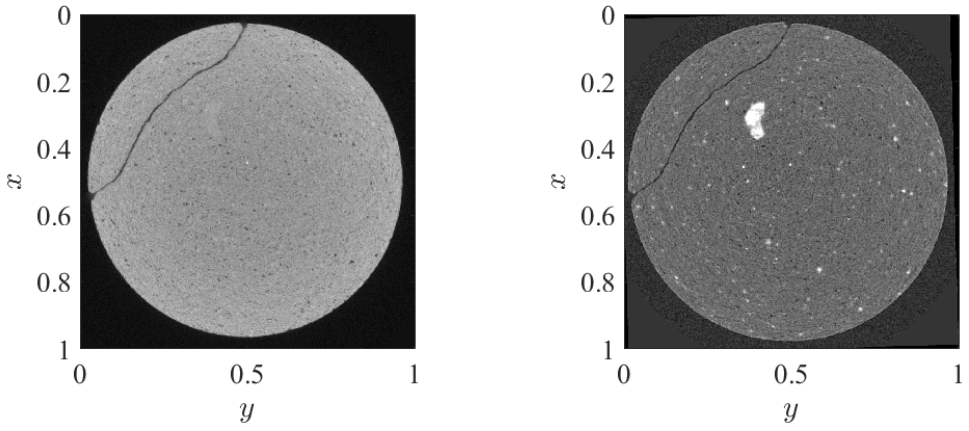

(a) X-ray (b) Neutron

Figure 1. An example pair of images obtained from slices manually selected from reconstructed a) x-ray and b) neutron tomography volumes

Dual X-ray and neutron tomographies as in Ref. [1], be they sequential or simultaneous, are very previous to study the mechanical deformation as they reveal different contrasts separate acquisitions and analysis where two DVC procedures are performed in parallel to study the mechanical deformation. Image registration renders possible the full use of these two sources of information as they can be related to a unique frame of reference.

There exist simple ways to find a transformation between two images of the same object acquired with different modalities. For example, a set of "landmark" points can be chosen in each image and, from their respective coordinates, the transformation $\boldsymbol{F}$ that registers these points and, thus, the two images can be determined [2]. To perform such an operation, a minimum of three points is needed in three dimensions, whereas two points are sufficient in two dimensions. Redundancy provided by using more points

allows robustness to be improved and, more importantly, to validate the transformation. This technique is expected to provide a matching with an accuracy of typically a few pixels. Moreover, the selection of landmark points may be tedious, especially in three dimensions.

Multimodal (multi-sensor) registration poses a particular difficulty due to the fact that the image contrast differs between different modality images. It is, therefore difficult to define a proper metric to assess the coincidence of the images. For this reason, artificial examples are often considered, for which the exact matching solution is known and quadratic differences (or equivalent signal to no ise ratio) between the test case and its known reference enables a given registration strategy to be assessed before application to real data [3]. The question of the representativeness of the artificial reference case and its noise remain mostly unknown. As discussed previously, multimodal registration can be achieved by selecting some features that are deemed to be more robust, such as sharp gradients, and, more specifically, their orientation, as the gradient magnitude cannot generally be compared. Strategies based on matching such features and attributing "points" for each match to maximize the registering score have been proposed [4], although they call for a necessarily subjective appreciation, as for the above mentioned landmark points.

Traditionally, methods based on mutual information (MI) [5] represent the leading technique in multimodal registration [6]. Put simply, the MI of two images is the amount of information that one image contains about the other and *vice versa* [7]. Such a metric provides a very general (probabilistic) framework that does not call for prior expectation or assumptions (such as linear relationships), but only requires a predictable relationship. However, one of its biggest pitfalls is the arbitrariness of the measure attached to "gray levels" and the lack of "distances" in the space of joint gray levels. Similarly, the knowledge of gray levels in the spatial vicinity of a point is not used. Another drawback is that most implementations require that the images to be registered are defined as random variables (*i.e.,* to determine their probability density functions), for which either discrete or continuous approaches need to be used (such as estimation through Parzen windows [8]).

The first approach for MI-based methods was proposed by Viola and Wells [8]. Subsequently, different approaches have been explored to obtain non-rigid multimodal registration: e.g., by using multiple local windows connected through a Gaussian window function ensuring continuity and smoothness of the deformation field [9]; by using a "correlation ratio" [10] based on the MI, but with consideration of the spatial information, which drastically improves the result; by combining MI with a term based on the image gradient to be registered [7] (multimodal images can have drastically different intensities, but, as the images fundamentally depict the same microstructures, the gradients in two multimodal images can be assumed to be similar); by modeling the deformed image as a viscous fluid that deforms under the influence of forces derived from the gradient of the MI [11] (built upon [12]).

Image registration and, in particular, Digital Image Correlation (DIC), or

specifically for 3D images, Digital Volume Correlation (DVC), has shown a much higher resolution potential for image matching [13]. Resolutions of one tenth to one hundredth of pixel can be achieved due to the exploitation of all pixels (or voxels) in a given region of interest. Such resolution is precious for solid mechanics applications [14]; hence, it would be appealing to benefit from a similar matching accuracy when dealing with images issued from different sensors. However, the difficulty is that a suitable metric has to be considered to estimate the "distance" between the two images when in coincidence.

Although not very frequent within the methodology of DIC, some prior work has explored the question of image registration with different modalities, but with a rather straightforward correspondence between the different contrasts. For instance, the calibration of distortions in an SEM has been studied aimed at registering a design of a speckle marking (a binary image) with an actual acquisition of the surface where the marking has been applied by a Pt e-lithography technique [15]. In this case an affine transformation of the gray levels was revealed to be sufficient. In Refs. [16, 17], the authors used the joint histogram of two simultaneous X-ray and neutron acquisitions assuming that the images were located identically without displacement and distortion. A more difficult example deals with an extension of classical stereo-vision or stereocorrelation [18, 13] to hybrid stereocorrelation. In this case, two images, one acquired with a standard optical camera and a second one with an IR camera, taken from two different points of view were matched so as to extract a 3D shape and displacement of the surface of a specimen [19].

The outline of the article is the following. After having introduced the challenge and notations in Section 2, it is proposed to address the registration as a minimization based on a "potential" Φ. The algorithm, including a multiscale version, is presented in Section 3 and follows closely approaches employed in global DVC, and for this reason, it will be referred to as multimodal DVC. Different formulations can be proposed for the potential as discussed in Section 4. However, within a probabilistic framework, the minimization of the potential can be seen as the maximization of a log-likelihood, which provides clues to formulate different descriptions. Two particular cases are proposed. The first is based on a Gaussian Mixture model to account for the joint histogram. This approach is shown to be very well suited to the problem at hand and has the additional benefit of leading to a segmentation of the images into phases. The second approach is model-free and, hence, does not rely on any specific prior modeling of the microstructure into phases. In contrast, the aim of this method is to learn the potential from the current joint histogram of gray levels. To illustrate the merits and limits of the proposed algorithm, Section 5 provides a demonstration of the application of the two approaches to a 3D test case. The results show that both approaches are successful at registering the images and provide comparable estimates of transformations. A brief summary of the proposed procedure and achieved results is given in Section 6 together with some perspectives.

2. Formulation

The two images to be matched are $f(\boldsymbol{x})$ and $g(\boldsymbol{x})$, where f and g denote the gray level values provided by each modality. The registration consists in finding the transformation $\boldsymbol{F}$ such that the material point at position $\boldsymbol{x}$ in the first image coincides with the same material point at position $\boldsymbol{x}' = \boldsymbol{F}.\boldsymbol{x}$ in the second one. To account for translation, rotation and dilation, it is convenient to use homogeneous coordinates $\boldsymbol{x} = (x_1, x_2, x_3, 1)$, where the last unit coordinate allows the change in origin to be considered within the same linear framework. The transformation is linear and, hence, can be represented by the 4×4 matrix $\boldsymbol{F}$ with a trivial conventional row $F(4,.) = (0,0,0,1)$. The non-trivial part is a 3×4 matrix, which may also be used to describe an arbitrary homogeneous deformation. The above assumptions imply that the deviatoric part of the strain (5 components) is null, so that the remaining degrees of freedom are $12 - 5 = 7$ consistent with the 3 translations, 3 rotations and 1 dilation. $\boldsymbol{F}$ can be restricted to the above 7 unknowns or left free with 12 unknowns, in which case the estimated deviatoric strain may be a way to validate the transformation or to detect possible distortions in one modality or the other.

In the following, the second possibility of 12 unknowns is chosen. More specifically, $\boldsymbol{F}$ is written as

$$\boldsymbol{F} = \begin{bmatrix} \begin{bmatrix} & & \\ & \boldsymbol{I} + \boldsymbol{\varepsilon} + \boldsymbol{\omega} & \\ & & \end{bmatrix} & \begin{Bmatrix} \\ \boldsymbol{t} \\ \\ \end{Bmatrix} \\ \{0 \quad 0 \quad 0\} & 1 \end{bmatrix}, \tag{1}$$

where $\boldsymbol{I}$ is the 3×3 identity matrix, $\boldsymbol{\varepsilon}$ and $\boldsymbol{\omega}$ respectively the symmetric and anti-symmetric parts of the $\boldsymbol{F}(1-3, 1-3)$ upper block, $\boldsymbol{t}$ is a translation vector. The origin of the coordinate system is located in one corner of the images (minimum x_i coordinates are 0). The anti-symmetric tensor $\boldsymbol{\omega}$ can be interpreted as a rotation, while $\boldsymbol{\varepsilon}$ is the infinitesimal strain tensor.

If both images were acquired with the same modality, it would be expected that for all $\boldsymbol{x}$ in the Region of Interest

$$f(\boldsymbol{x}) = g(\boldsymbol{F} \cdot \boldsymbol{x}), \tag{2}$$

i.e., the brightness conservation assumption, which is the basis of digital image correlation [14]. The above relationship can be relaxed, for instance to account for different brightness or contrast between the two images, and rewritten as

$$\Phi(f(\boldsymbol{x}), g(\boldsymbol{F} \cdot \boldsymbol{x})) = 0, \tag{3}$$

where $\Phi(f, g) = \alpha_1 f - \alpha_2 g - \alpha_3$ is an affine transformation. Image correlation based on normalized cross-correlation, for instance, is naturally immunized against such an affine transform Φ [13]. Image correlation is also generally formalized to tolerate a gaussian white noise in both images. This leads to a relaxation of the above formulation so

that the transformation $\boldsymbol{F}$ is given by the minimizer of the functional $\mathcal{T}$ over all affine transforms,

$$\begin{aligned} \boldsymbol{F} &= \text{Argmin}_{\boldsymbol{G}} \mathcal{T}[\boldsymbol{G}], \\ \mathcal{T}[\boldsymbol{G}] &= \sum_{\boldsymbol{x} \in ROI} \Phi^2(f(\boldsymbol{x}), g(\boldsymbol{G} \cdot \boldsymbol{x})) \ . \end{aligned} \tag{4}$$

When correlating images generated by different modalities, the change in gray levels is generally much more drastic than just an affine transformation. Yet, a physically sound assumption can be made that the material consists of few phases $i = 1, ..., N$ that have well defined signatures (f_i, g_i) in the (f, g) plane. In other words, wherever a phase is located in space, it will give the same gray level within the same modality image. Note however that, even if the above statement seems quite natural, for instance the artifact called "beam hardening" in X-ray tomography [20] violates this assumption. Figure 2 shows a log-scale joint histogram of the two images shown in Figure 1 prior to any registration. It is natural to interpret the most salient peaks as corresponding to the different distinct phases (f_i, g_i), with the reservation that, prior to registration, a wrong voxel positioning may give rise to spurious peaks, as the f and g coordinates may refer to a different phase.

The spirit of the proposed approach is to design a "potential" Φ, which is locally minimum for each of the different phases (f_i, g_i), and to determine $\boldsymbol{F}$ from the minimization of $\mathcal{T}$. Different approaches can be followed to choose this potential, the discussion of which is deferred to a later section. The minimization algorithm is first discussed in the following.

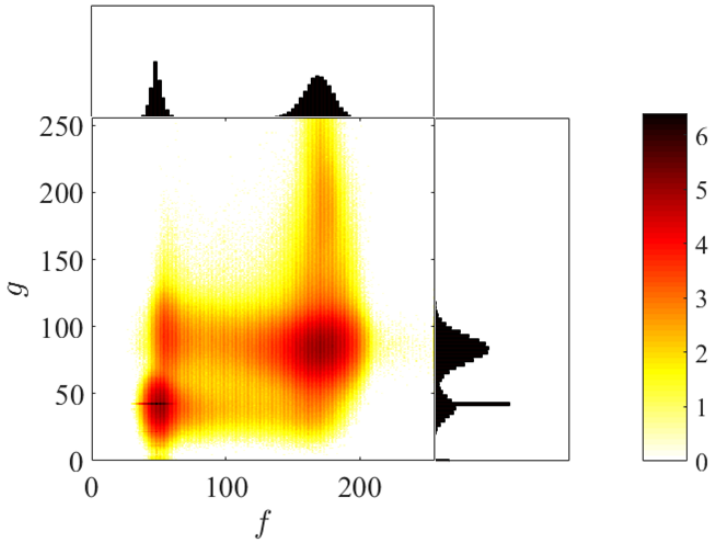

Figure 2. $\log_{10}$-scale joint histogram of gray levels from the sections shown in Fig. 1 with no registration. $f(\boldsymbol{x})$ is the x-ray image, and $g(\boldsymbol{x})$ is the neutron image. Histograms of each individual image are presented in black on a linear scale.

3. Registration algorithm

The registration of two images resulting from different modalities is a nonlinear problem and it is proposed to solve this with a Gauss-Newton algorithm, with the assumption that the initial determination of $\boldsymbol{F}$ is close to the solution. Thus, the algorithm consists in the determination of successive corrections to the current estimate of the transformation, $\boldsymbol{F}^{(n)}$, at iteration n.

3.1. Gauss-Newton algorithm

The following notations are introduced for the corrected image

$$\widetilde{g}^{(n)}(\boldsymbol{x}) \equiv g(\boldsymbol{F}^{(n)} \cdot \boldsymbol{x}). \tag{5}$$

The correction at each iteration, n, is denoted as $\boldsymbol{G}^{(n+1)} = \boldsymbol{I} + \boldsymbol{\delta F}^{(n+1)}$, such that $\boldsymbol{F}^{(n+1)} = \boldsymbol{G}^{(n+1)}\boldsymbol{F}^{(n)}$. The functional to be minimized takes the following expression

$$\mathcal{T}[\boldsymbol{F}^{(n+1)}] = \sum_{\boldsymbol{x}\in ROI} \Phi^2(f(\boldsymbol{x}), \widetilde{g}^{(n)}(\boldsymbol{G}^{(n+1)} \cdot \boldsymbol{x})). \tag{6}$$

Expanding the above integral up to second order in $\boldsymbol{\delta F}^{(n+1)}$ yields

$$\begin{aligned} \mathcal{T}[\boldsymbol{F}^{(n+1)}] &= \sum_{\boldsymbol{x}\in ROI} \Phi^2\left(f, \widetilde{g}^{(n)} + (\boldsymbol{\nabla}\widetilde{g}^{(n)} \cdot \boldsymbol{\delta F}^{(n+1)} \cdot \boldsymbol{x})\right) \\ , &= \sum_{\boldsymbol{x}\in ROI} \begin{array}{l} \left[[\Phi^2] + ([\Phi^2]_{,2})(\boldsymbol{\nabla}\widetilde{g}^{(n)} \cdot \boldsymbol{\delta F}^{(n+1)} \cdot \boldsymbol{x})\right. \\ \left. + (1/2)([\Phi^2]_{,22})(\boldsymbol{\nabla}\widetilde{g}^{(n)} \cdot \boldsymbol{\delta F}^{(n+1)} \cdot \boldsymbol{x})^2\right] \end{array} , \end{aligned} \tag{7}$$

where the notations $[\Phi^2]_{,2}$ and $[\Phi^2]_{,22}$ are used to refer to the partial derivatives with respect to the second argument once or twice, respectively, and Φ^2, as well as its derivatives, are computed at point $(f, \widetilde{g}^{(n)})$ (omitted for conciseness).

Because the unknowns are themselves the components of a matrix, $\delta F_{ij}^{(n+1)}$, it is convenient to relabel them with a single index m such that $\delta F_{m}^{(n+1)} \equiv \delta F_{i(m)j(m)}^{(n+1)}$, thus defining implicitly $i(m)$ and $j(m)$, in the spirit of Voigt notation for the strain tensor [21]. Since the above expression of the functional is quadratic in $\boldsymbol{\delta F}^{(n+1)}$, its determination is the solution to a linear system

$$M_{mp}^{(n)} \delta F_{p}^{(n+1)} = A_{m}^{(n)}, \tag{8}$$

where

$$M_{mp}^{(n)} = \sum_{\boldsymbol{x}\in ROI} ([\Phi^2]_{,22})(x_{j(m)}\widetilde{g}_{,i(m)}^{(n)})(x_{j(p)}\widetilde{g}_{,i(p)}^{(n)}) \tag{9}$$

and

$$A_{m}^{(n)} = -\sum_{\boldsymbol{x}\in ROI} ([\Phi^2]_{,2})(x_{j(m)}\widetilde{g}_{,i(m)}^{(n)}). \tag{10}$$

In the preceding equations, x_i refers to the i-th component of the vector $\boldsymbol{x}$.

The correction to the transformation at each iteration is taken into account as

$$\boldsymbol{F}^{(n+1)} = (\boldsymbol{I} + \boldsymbol{\delta F}^{(n+1)})\boldsymbol{F}^{(n)}. \tag{11}$$

It may be worth noting that standard DIC and DVC (*i.e.,* with two images derived from the same modality) is recovered for $\Phi^2(f,g) = (1/2)(f-g)^2$, where the Hessian and second member take, respectively, the well-known [22], following expressions,

$$M_{mp} = \sum_{\boldsymbol{x}\in ROI} (x_{j(m)}\widetilde{g}^{(n)}_{,i(m)})(x_{j(p)}\widetilde{g}^{(n)}_{,i(p)}) \tag{12}$$

and

$$A_m = \sum_{\boldsymbol{x}\in ROI} \left(f - \widetilde{g}^{(n)}\right)(x_{j(m)}\widetilde{g}^{(n)}_{,i(m)}). \tag{13}$$

Within the classical DIC and DVC frameworks, a very useful tool is the residual field, $\rho(\boldsymbol{x})$, namely the image difference between reference and corrected deformed images. This field, which is resolved at the voxel scale, indicates the success of the registration, spatially over the image space, and, in particular, where the registration is not successful. Based on this field, the registration may be revisited with other assumptions or corrections. When extended and generalised to the case of intermodal registration, it is observed that the equivalent residual field, computed at convergence, is

$$\rho(\boldsymbol{x}) = \frac{\partial \Phi^2(f(\boldsymbol{x}), g(\boldsymbol{F}\cdot\boldsymbol{x}))}{\partial g}. \tag{14}$$

3.2. Multiscale approach

One of the major difficulties encountered in DIC and DVC is the initialization of the transformation $\boldsymbol{F}$. If the initial solution is too far from the actual solution, the proposed algorithm may simply diverge, or converge to a secondary minimum of $\mathcal{T}$. In this context, convergence is assessed based on the ratio of the maximum displacement over all voxels in the region of interest and the correlation length of the image, as estimated from the pair correlation function. When this ratio exceeds unity, there is a risk of convergence toward spurious secondary minima.

One easy (and cheap) way of enhancing the robustness of DIC and DVC is to coarse-grain the images [23, 24]. This coarse-graining consists of a low-pass filtering with a filter at a characteristic scale of 2 pixels, and a resampling that retains only one pixel out of two in each spatial direction. The result is two smaller images whose sizes have been reduced in size by a factor of 2^d, where d is the space dimension. The resulting pixels can be seen as "super-pixels", which account for the underlying group of 2^d fine scale pixels. The simplest coarse graining step consists in replacing each elementary group of 2^d pixels, by a single super-pixel with the arithmetic average of the gray levels of the group. This coarse-graining operation can be applied recursively, thereby defining a pyramidal construction, where each pyramid level represents the coarsened image of the level below. When expressed in super-pixels, the maximum displacement decreases by a factor of 2 at each level of the pyramid. Conversely, the correlation length cannot decrease below one pixel, and it is, therefore, easy in a few pyramidal levels to restore a safe convergence criterion. Once convergence has been

achieved at a given level n of the pyramid, the resultant transformation may be used to initialize a new correlation computation at the level $(n-1)$, down to level 0 which conventionally is the original image resolution. In the implementation discussed below, a Gaussian pyramid implementation [25] has been chosen.

In the particular context of multimodal DVC, the only new question to address is how to design the potential Φ at different levels of the pyramid. Pure phases that are present as large, relatively homogeneous regions tend to display sharper peaks in the joint histogram as the pyramid level increases and noise is suppressed (by the averaging in the coarsening operation). Phases that occur only in very small regions are progressively assimilated with their surroundings and contribute to broadening of their distribution in the joint histogram. Therefore, there is no general rule as to how to define the potential, Φ, that would be independent of the material microstructure. Consequently, it is suggested to adjust the Φ at each level of the pyramid, based on the respective joint histogram. The adopted multiscale procedure is summarized in Algorithm 1.

Algorithm 1 Multiscale intermodal DVC algorithm

```
Initialize F
for PyramidLevel=n:-1:0 do
    Compute pyramidal f(x)
    Compute pyramidal g(x)
    Compute g̃^(0)(x)
    Compute joint histogram (f, g̃^(0))
    Adjust Φ
    while ‖δF‖ > ε do
        Compute M_mp
        Compute A_m
        Solve M_mp δF_p = A_m
        Update F
        Compute g̃^(n)(x)
    end while
end for
Display residual "distance" (Φ(f, g̃^(n))^2)_,2
Compute joint histogram
Extract segmentation
```

4. Design of potential Φ

4.1. What is expected from the potential?

In a simple case of multimodal registration of images containing only a few pure phases, it is natural to tailor a specific $\Phi(f,g)$ functional relying on very few parameters. One

practical way to design the potential, in such a case, is to consider the joint histogram of gray levels and select the points that correspond to the centroids of the data clusters (which are assumed to be related to the occurrence of a particular phase). Figure 2 shows joint histogram for the sections shown in Fig. 1 that has well defined data clusters, suggesting the existence of well-defined phases. It is to be noted that when images are not yet registered, the joint histogram will display spurious populations of voxels whose gray values (f, g) correspond to different phases, as mentioned earlier. Those points share the abscissa and ordinate of existing (but different) phases. It will be shown in the following that the region around $f \approx 55$ and $g \approx 95$ in Fig. 2 is such a spurious point.

Pixel (or voxel) data are not independent data with respect to their neighbours. In reality they rather correspond to an integration over a finite area or volume. Thus, voxels are not expected to consist only of pure phases. For those voxels lying across the boundary between two phases i and j, it is expected that the voxel value will be $f = \alpha f_i + (1 - \alpha) f_j$ and $g = \alpha g_i + (1 - \alpha) g_j$, where α ranges continuously over the interval $[0, 1]$. Hence a distribution of gray levels along the segment connecting two pure phases is expected. The probability density along this line is related to the density of i-j grain boundaries crossed by voxels. The probability is low for a coarse microstructure and increases for a finer one. Consequently, a term that lowers the penalty given to a mixture between phases may be introduced in the potential Φ.

It should be emphasized that the image resolution itself may have a very significant impact on the joint histogram and hence on the potential to be used. Although trivial, this observation justifies that the potential should, in principle, be adjusted at each level of a multiscale procedure.

4.2. Likelihood interpretation

The previous discussion points to the design of specific Φ functionals, based on the observation of joint density distribution for $\boldsymbol{h} = (f, g)$ gray levels. The latter potential should ideally be adjusted once registration has been achieved, and this is the motivation for computing Φ^2. At best, an iterative treatment of the problem may be envisioned where improved approximations of the potential are defined at each iteration.

The previous approach can be rewritten in terms of a probabilistic inference approach. From a registered pair of images, one may compute the probability distribution function of $\boldsymbol{h}$, $p(\boldsymbol{h})$. Thus, the likelihood that the two images $f(\boldsymbol{x})$ and $g(\boldsymbol{F} \cdot \boldsymbol{x})$ coincide can be written

$$\mathcal{L}[\boldsymbol{F}] \propto \prod_{\boldsymbol{x} \in ROI} p(f(\boldsymbol{x}), g(\boldsymbol{F} \cdot \boldsymbol{x})). \tag{15}$$

The log-likelihood assumes a convenient form,

$$\log(\mathcal{L}[\boldsymbol{F}]) = \sum_{\boldsymbol{x} \in ROI} \log(p(f(\boldsymbol{x}), g(\boldsymbol{F}.\boldsymbol{x}))) + \text{constant}. \tag{16}$$

Maximizing the probability or equivalently log-likelihood should, thus lead to the registration solution. Hence, it is deduced from the comparison with the postulated form of the DVC functional $\mathcal{T}$ that a natural expression of the potential, Φ^2, is (disregarding an unnecessary additional constant),

$$\Phi^2(\boldsymbol{h}) = -\log(p(\boldsymbol{h})). \tag{17}$$

Rather than Φ^2 being a subjectively designed potential, the above log-likelihood interpretation provides a well-defined and objective foundation. The difficulty however, is that the joint distribution and hence the cost function can only be known once registration has been achieved, and the latter calls for the cost function.

A first strategy can be designed based on the assumption that the starting point of the registration is already good enough to provide a reasonable estimate of Φ^2. From this initial estimate, a registration transformation $\boldsymbol{F}$ is estimated and a corrected joint histogram is computed, from which Φ^2 can be updated. Subsequently a fixed point solution, correspond to registration, is sought by iteration. To avoid the solution becoming trapped in local minima and to compute gradients and curvatures safely, the joint histogram p is first filtered by convolution with a Gaussian kernel whose standard deviation is 3 gray levels (on an 8-bit scale), which is then used to compute Φ^2 as in Eq. 17. This approach is denoted as LL (for log-likelihood) in the following.

In contrast to several of the approaches cited in the introduction, the above procedure for determining Φ is no longer just a reasonable choice, or a particular heuristics, such as the quadratic differences, peak SNR or mutual information. In fact, the procedure exploits a functional that provides a maximum likelihood estimator without any *a priori* assumption. The only caveat is that no guarantee for convergence can be shown. It is however clear that close to the solution a converged solution that coincides with the desired one can easily be obtained. The principal limitation is, therefore, the maximum tolerable distance from the solution that still leads to the solution. In this respect, the multiscale procedure makes the algorithm more robust to large initial mismatches.

4.3. Gaussian mixture approach

A classical way to describe a two-parameter distribution is the Gaussian mixture model [16, 17], which consists in representing the distribution of gray levels pairs $\boldsymbol{h} = (f, g)$ as

$$p(\boldsymbol{h}) = \sum_i \phi_i \mathcal{N}(\boldsymbol{h}; \boldsymbol{\mu}_i, \boldsymbol{\sigma}_i), \tag{18}$$

where $\mathcal{N}(\boldsymbol{h}; \boldsymbol{\mu}, \boldsymbol{\sigma})$ is the normal distribution of $\boldsymbol{h}$ with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\sigma}$. Introducing

$$\lambda_i(\boldsymbol{h}) = (1/2)(\boldsymbol{h} - \boldsymbol{\mu}_i)\boldsymbol{\sigma}_i^{-1}(\boldsymbol{h} - \boldsymbol{\mu}_i), \tag{19}$$

p reads

$$p(\boldsymbol{h}) = \sum_i \exp(-\lambda_i(\boldsymbol{h}) + \log(\phi_i)). \tag{20}$$

If the Gaussians are well separated, then

$$\begin{aligned} p(\boldsymbol{h}) &\approx \max_i \left[\exp(-\lambda_i(\boldsymbol{h}) + \log(\phi_i))\right] \\ &\approx \exp\left[-\min_i(\lambda_i(\boldsymbol{h}) - \log(\phi_i))\right]. \end{aligned} \tag{21}$$

Using the previous relationship between p and Φ^2, the following simple expression of the potential results

$$\Phi^2(\boldsymbol{h}) = \min_i(\lambda_i(\boldsymbol{h}) - \log(\phi_i)). \tag{22}$$

This second strategy (labeled GM in the following) appears comparable to the previous one in the sense that Φ^2 still represents the negative log-likelihood. However, as compared to the previous LL approach, which is a plain depiction of the joint histogram, a specific model for the joint distribution is proposed that naturally focusses on the most populated peaks. Therefore spurious secondary peaks may be ignored if the number of Gaussian components is small enough and if an approximate registration is available initially. Thus, this GM approach is expected to help convergence from more distant initial registration than LL. However, because the approximation involved forces a simplifying model in GM, it is anticipated that the uncertainty of the finally obtained registration will be larger than that of the LL approach.

A side benefit of the GM approach is that, as the potential appears as consisting of a few paraboloids, the $\boldsymbol{h}$ plane is naturally partitioned into sub-domains each belonging to a single paraboloid; the boundaries between subdomains being simple conic curves. The label i where the minimum is reached in Eq. 22 can be assigned to each gray level pair $\boldsymbol{h}$, $\psi(\boldsymbol{h}) = i$. Each subdomain i is characterized by a "phase" whose signature is the gray level pair $\boldsymbol{\mu}_i$ of the apex of the paraboloid. Moreover, a linear function and a constant function can be introduced, respectively, for $[\Phi^2]_{,2}$ and $[\Phi^2]_{,22}$, in each subdomain of this partition. After registration, assigning the label $\psi(\boldsymbol{h})$ to each voxel where the gray level pair is $\boldsymbol{h}$ allows the image to be directly segmented based on the signatures of each phase. Such a segmentation is not directly available following the LL approach, although one could still consider the basin of attraction of each minimum of the Φ^2 potential using a steepest descent.

5. Test case

To illustrate the performance of the two approaches proposed above, LL and GM, a 3D case is chosen where two imaging modalities, X-ray and neutron tomography, have been used to image the same sample with similar resolutions.

5.1. Experimental details

The studied material is a Bentheim sandstone with an average porosity of about 23 %, a mean grain diameter of 300 μm and a composition of 95 % quartz, 3% kaolinite and 2 % orthoclase [26]. The sample shape is cylindrical with 50 mm diameter and 100 mm

height. Tomographic imaging of the specimen was made independently using X-ray and neutron tomography.

Both the x-ray and neutron measurements were carried out at the Helmholtz-Zentrum Berlin (HZB). Neutron images were acquired at the beamline CONRAD [27] with L/D of 500 (where D is the pinhole aperture defining the neutron beam size and L is the distance from the pinhole to the sample). The acquisition time for each radiography was 30 seconds, which gives a total time of about 6 hours for a complete scan of 600 images. The resulting reconstructed tomography volume images had a voxel size of about 30 μm. X-ray tomographies were acquired using a lab source running at a voltage of 120 kV, a current of 83 μA giving a spot size of about 10μm. The acquisition time for each radiography is 1.3 s. 1300 different angles are acquired spread through 360°, with 5 radiographies averaged at each angular position, giving a total scanning time of about 2.5 hours. The resulting reconstructed tomography volume images had a voxel size of about 30 μm.

In the following, a Matlab implementation has been used to illustrate the feasibility of the approach (this implementation has not yet been optimised in terms of computational efficiency). To limit memory usage, two volumes of height $\Delta z = 128$ voxels where first extracted from both 3D-images. These subvolumes were then 2×2-binned in the transverse (x, y) directions, so that the cross-sections were 900×900 voxels for both modalities, and re-sampled into 8 bits. A four-level pyramidal scheme was used where level 1 corresponds to the $900\times900\times128$ volume. However, to preserve sufficient thickness at all levels of the pyramid, it was chosen not to coarse-grain the z direction. Hence at level 4, the volume shape was approximately cubic $112\times112\times128$, but the microstructure in the images was highly anisotropic, as each level-4 voxel resulted from a $8\times8\times1$ group of voxels in the original image. The price to pay for this anisometric re-scaling is that one should be careful when transferring the result of one level to be the initialization of the next one. However, resorting to physical coordinates, rather than voxel numberings, is a simple way to avoid ambiguities.

5.2. Results of the GM approach

The first test was performed using the Gaussian mixture (GM) approach developed in Section 4.3. Although automatic selection algorithms are available to fit Gaussian mixture distributions, our attempts to use them lead to rather poor fits (more often than not the means were pushed to the limits of accessible gray levels, 0 or 255). Therefore, a simple procedure was designed for fitting the highest peaks, which involved fitting the highest of the residual between the measured joint histogram and the already identified Gaussian components recursively. Albeit simple, this procedure was found to be robust and adequate for our usage. The number of components was *a priori* prescribed from the shape of the joint histogram. The latter was selected independently per pyramid level, as it is highly dependent on the level. Here, this shape went from two components for level 4 to four components at level 1. Choosing two Gaussian components at the

coarsest level forbids the spurious third peak (see Fig. 3(a)) to be mistaken for a real one, thereby favoring its annihilation with the transformation $\boldsymbol{F}$.

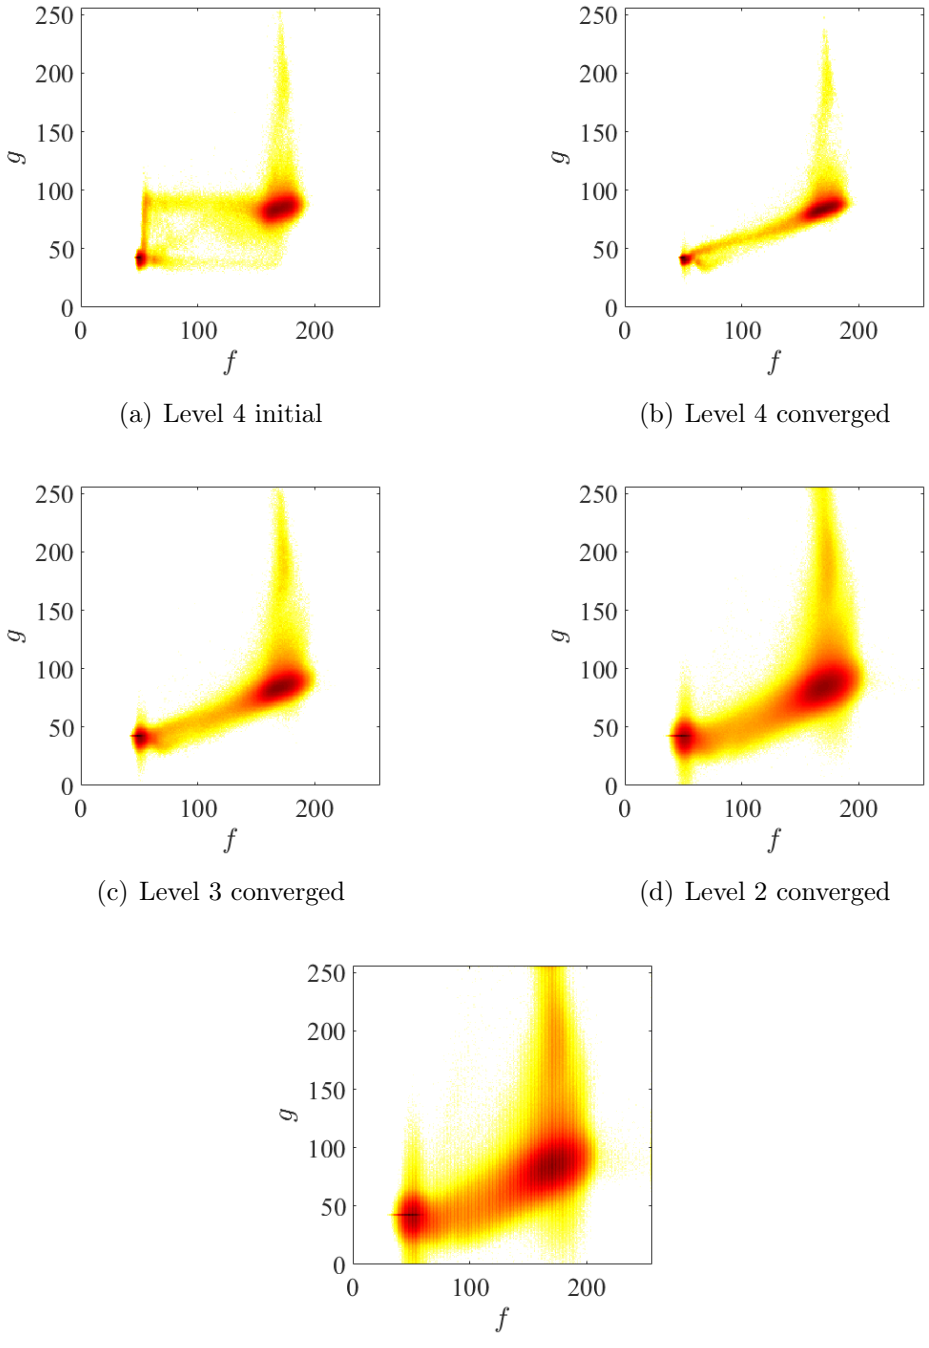

(a) Level 4 initial

(b) Level 4 converged

(c) Level 3 converged

(d) Level 2 converged

(e) Level 1 converged

Figure 3. (a) $\log_{10}$-scale joint histogram at various stages of the GM registration from (a) the initial step at level 4, and after registration at levels 4 to 1 respectively in (b) to (e) of registered image

Figures 3(b-e) show the change of the joint histogram of both images after

registration at different levels of the pyramid. It is observed that the spurious patterns that appear as horizontal or vertical lines in Fig. 3(a) already disappeared after convergence at level 4. The joint histogram appears to display very few signs of mismatch already at this stage. At lower levels, only a broadening of the histograms is visible, which simply results from the higher resolution and the higher noise magnitude.

Table 1. Evaluation of the transformation parameters at convergence for each pyramid level, using the GM approach. The convergence criterion was chosen to be such that the norm of the deformation gradient part of $\boldsymbol{\delta F}$ should be less than 5×10^{-4}

Pyramid level (iterations)		$L=4$ (47)	$L=3$ (12)	$L=2$ (12)	$L=1$ (10)
ϵ_{xx}	(%)	1.78	1.72	1.69	1.66
ϵ_{yy}	(%)	1.75	1.68	1.63	1.60
ϵ_{zz}	(%)	-2.04	-2.90	-2.72	-2.88
ϵ_{yz}	(%)	-0.36	-0.24	-0.21	-0.21
ϵ_{xz}	(%)	-0.21	-0.01	0.03	0.03
ϵ_{xy}	(%)	-0.05	0.01	0.04	0.04
ω_x	(deg.)	-0.11	-0.07	-0.06	-0.07
ω_y	(deg.)	0.43	0.52	0.58	0.59
ω_z	(deg.)	0.59	0.68	0.70	0.71
T_x	(vox.)	-14.1	-15.7	-16.1	-16.0
T_y	(vox.)	-11.8	-10.1	-9.4	-9.1
T_z	(vox.)	9.8	9.1	9.3	9.6

Table 1 gives the number of iterations at each level of the pyramid, together with the estimated parameters of the transformation $\boldsymbol{F}$. The convergence criterion is chosen to be the stationarity of the deformation gradient, and relaxation is stopped when the change between two iterations is less than 5×10^{-4} or when the number of iterations reaches a maximum number, here chosen to be 60 (with a minimum of 10). It is observed that the number of iterations is large at level $L = 4$ (where the volume is much smaller than at later stages and hence computation is much faster). After this first step, the number of iterations is close to the imposed minimum. Consistently, it is observed that the transformation is already very well determined after convergence at level $L = 4$. Further corrections from one level to the next are of order 10^{-3} in strain. Only the ϵ_{zz} strain shows a larger correction from $L = 4$ to $L = 3$, where it almost reaches its final stabilized value. It is also interesting to observe that directions x and y behave in a similar fashion, with very close values. However, they both differ quite significantly from direction z. Since the latter is the rotation axis, it is deduced that it is not coupled with the perpendicular directions that, in contrast, are processed together in the reconstruction and, hence, cannot be different. One possible explanation of this difference may result from an imperfect alignment of specimen rotation axis or camera.

As above discussed, a residual field, $\Phi^2_{,2}$, is constructed at convergence; see Figure 4(a). The fact that neither f or g, can be recognised in the residual field is an indication that the result can be trusted.

As discussed earlier, an additional output of the proposed registration is a natural segmentation of the registered images and labelling of each voxel as a specific "phase". This segmentation is achieved by starting from any gray level pair (f, g) and following the steepest descent in the potential, which directly provides a partition into (here) four basins of attraction. By this approach, each (f, g) pair can be labelled by the indices of the different basin roots. Figure 4(b) shows the result of the labelling issued from the GM approach for the test case. It is noteworthy that no other treatment has been performed (such as spatial filtering or any other morphological operator) to dampen the noise.

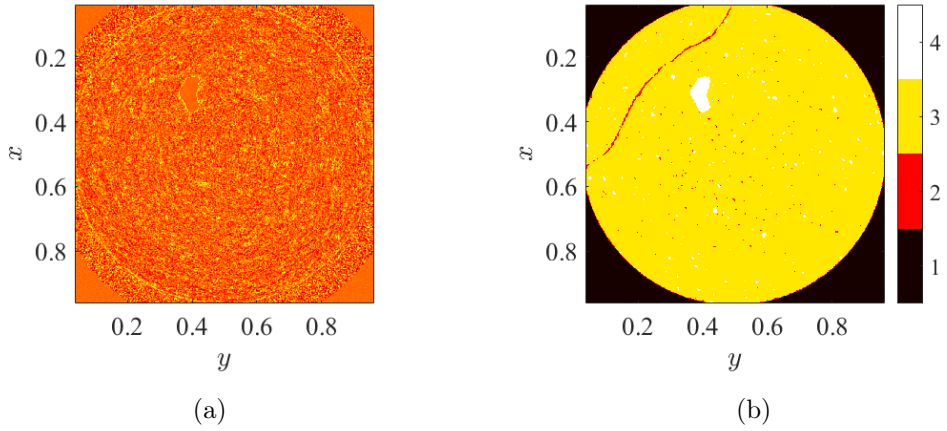

Figure 4. (a) Residual at convergence (level 1) and (b) segmented map (with four phases) after registration using the GM approach (section along median z plane)

5.3. Results of the arbitrary likelihood (LL) approach

The advantage of the LL approach, over the GM method, lies in the fact that there is less room for personal judgement, as the potential is based on the observed histogram. However, as could be seen in the previous section in Fig. 3, the joint histogram itself appears to change significantly depending on the pyramid level. In particular, registration is rather poor at the initial state and the resulting spurious features in the joint histogram will not be corrected in the LL approach as they were in the GM approach (simply because very few (*i.e.,* 2) components were chosen initially). In the LL approach, the only processing of the joint histogram is a convolution by a Gaussian whose width is three gray levels to remove high frequency fluctuations and hence local minima.

Figure 5 shows the initial potential at different pyramid levels together with the converged joint histogram (using the same log-scale). As anticipated, because some

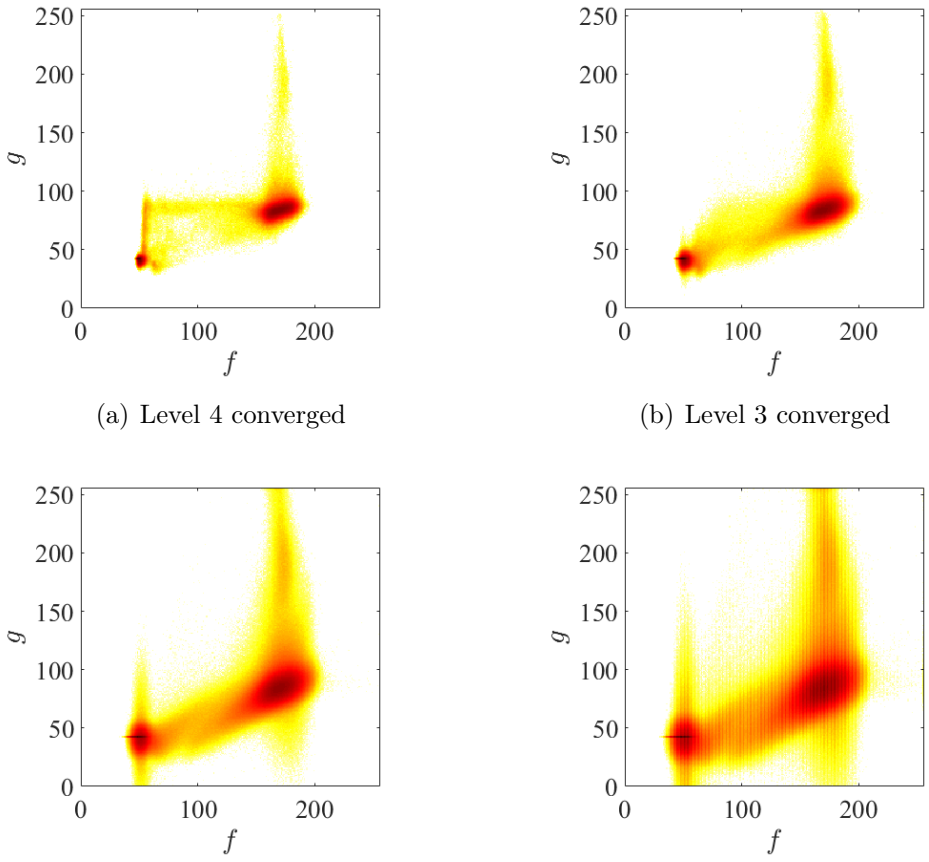

(a) Level 4 converged

(b) Level 3 converged

(c) Level 2 converged

(d) Level 1 converged

Figure 5. Observed joint histograms after convergence at four pyramidals level from $L = 4$ to $L = 1$

spurious features are present in the initial histogram, the early convergence is quite poor and at the end of the $L = 4$ relaxation, very few changes occur. In spite of this, at level $L = 3$, the vertical branch at lower f values disappears almost entirely. The following levels and athe final joint histogram are similar to the GM case.

Table 2 gives the results of the DVC registration based on the LL approach. It is observed, first, that the number of iterations in the higher levels is much larger than for the GM approach, as the maximum iteration number is reached. It is also observed, consistent with the previous observations, that the estimation of the transformation $\boldsymbol{F}$ at level 4 is totally wrong. After convergence at level $L = 3$, more acceptable values are obtained showing, for the most part, a good stability until the end of the registration at $L = 1$. In particular $L = 2$ and $L = 1$ results show only minor changes. Moreover, some of these estimates of the transformation $\boldsymbol{F}$ are in very good agreement with the previous approach. All strain components differ from the GM estimates by less than 10^{-4}, but the ϵ_{zz} whose difference with the GM estimate is close to 1%. Rotations are also almost identical. Translation differences are also in the voxel range (and 2 voxels

along z). One possible reason for larger differences in the z-direction may come from the fact that the sample size is also much shorter along the z axis (128) than in the transverse directions (900).

Table 2. Evaluation of the transformation parameters at convergence for each pyramid level, using the LL approach. The convergence criterion was chosen to be such that the norm of the deformation gradient part of $\boldsymbol{\delta F}$ should be less than 5×10^{-4}

Pyramid level (iterations)		$L=4$ (60)	$L=3$ (60)	$L=2$ (41)	$L=1$ (10)
ϵ_{xx}	(%)	0.24	1.58	1.63	1.60
ϵ_{yy}	(%)	0.10	1.39	1.53	1.52
ϵ_{zz}	(%)	-2.91	-3.87	-1.82	-1.96
ϵ_{yz}	(%)	-0.64	-1.38	-0.24	-0.21
ϵ_{xz}	(%)	-0.26	0.66	0.01	0.03
ϵ_{xy}	(%)	0.10	0.16	0.03	0.04
ω_x	(deg.)	-0.22	-0.61	-0.03	-0.02
ω_y	(deg.)	0.15	0.85	0.70	0.72
ω_z	(deg.)	0.13	0.66	0.73	0.73
T_x	(vox.)	-4.2	-16.7	-16.1	-15.9
T_y	(vox.)	0.7	-6.4	-7.9	-7.9
T_z	(vox.)	9.5	12.1	11.6	11.6

Figure 6(a) shows the residual map at the final step of the LL registration procedure for the case example. This residual map indicates a good convergence where no bias is visible, as for the GM approach. Also as for the GM approach, the different Gaussian components can be associated with a specific material phase at convergence. For the LL case, a similar, albeit less obvious procedure, to segment the phases is designed by identifying pure phases (*i.e.,* labels) at maxima of the joint distribution. Any gray level pair will flow, with a steepest ascent algorithm, toward a maximum that defines the corresponding phase. This procedure groups gray levels into domains from the basin of attraction of the maxima. The smoothing of the distribution using a Gaussian filter, as earlier discussed for the LL approach, limits the number of maxima. Figure 6(b) shows the distribution of phases, where the number of phases has been set to 6 without any prior subjective statement.

To check whether the observed difference in the transformations determined using the two proposed approaches is due to a poor convergence, initialization of the LL approach with the result from the GM one or vice versa was tried. Each approach led to the same results as those already reported (*i.e.,* with no influence of the initialization), so that slow convergence cannot be blamed for the difference. Thus, the observed difference in the transformation between the two approaches presumably comes from the joint distribution function, which is not adequately depicted in the GM approach.

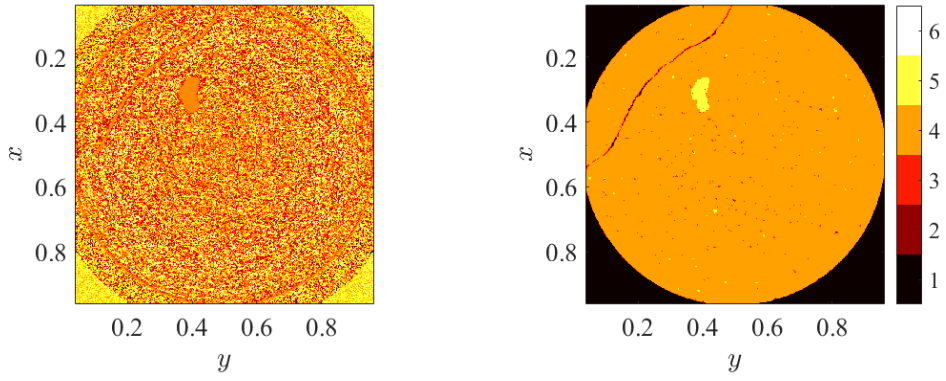

(a) Level 1 (b) Level 1

Figure 6. (a) Residual field after registration using the LL approach. (b) Segmented map (section along median z plane)

In contrast, the LL approach, for which priors are not postulated, appears to be more satisfactory, although it is difficult to provide a more objective appreciation of this from the registration results.

Finally, it is worth emphasizing that tailored potentials may be designed on demand to achieve a registration based on a specific prior judgement of the reliable region of the joint histogram. However, a systematic procedure that would allow for a simplified and yet faithful picture of the joint distribution is the most desirable. Two possible methods to describe the potential have been shown as examples, the first one, GM, was presumably oversimplified while the second, LL, did not provide enough constraint.

To check the quality of registration, Fig. 7 shows two complementary, composite mosaic pictures assembled from checkerboards of squares extracted from one modality or the other , after registration. These composite-mosaic images allow continuity to be seen with ease whenever relevant, and also to observe that some patterns are hardly visible in one modality, but very noticeable in the other one.

6. Conclusions

It has been shown that the registration of two images acquired with different modalities can be performed through an extension of Digital Volume Correlation through the minimization of a potential Φ^2. Furthermore, it was shown that, if this potential coincides with (the cologarithm of) the joint probability distribution, then the registration solution is the one that maximizes the likelihood. From this observation two variants for the potential were formulated:

- a Gaussian mixture model (GM) where the joint distribution is fitted to a set of a few Gaussian distributions
- a log-likelihood (LL) approach where the measured joint distribution is used directly

to design the potential for the registration, which is a less intrusive choice.

A real case study on two 3D volume images of the same test object, but derived from two different imaging modalities, namely x-ray and neutron tomographies, allowed the performance of the two above mentioned algorithms to be investigated. The GM algorithm was found to be very stable with a fast convergence in the multiscale framework. Although less stable and displaying a slower convergence, the LL algorithm was also shown to be able to provide a good registration, for which the mean transformation parameters were in good agreement with those obtained from the GM algorithm.

A voxel accuracy can be reached, revealing, in the case under study, an unexpected artifact, in the form of an apparent strain between the images corresponding a few percent difference of the dimensions parallel to the tomographic rotation axis direction compared to the dimensions perpendicular to this. One may speculate that this may result from a poor detector alignment.

Because the proposed procedure has been applied on the registration of two images of the exact same sample, a simple space regularization (with 12 degrees of freedom) has been chosen. However, in the case where the specimen is subjected to deformations between the two scans (as due to drying, swelling, creep, as a response to mechanical testing, or even for living specimen that may move, or breathe, ...), the degrees of freedom could be enriched as classically done with DVC.

An additional, "for free" result from the proposed registration method is a relatively straightforward segmentation of material phases according to their contrasts in both modalities. The registration of such two-modality images also opens the way to enrich/correct a low resolution, low contrast, or noisy image from one modality with the other one. This could be very beneficial, for example, to the exploit simultaneous, dual X-ray and neutron tomographies (as opposed to separate acquisitions as in the example presented herein from [1]) to profit from potentially faster X-ray tomography acquisition to enhance the neutron tomography and to extract greater information from the combination of the different sensitivities to the different phases of the two modalities (cf. project "NeXT-Grenoble" [28]).

Acknowledgement

The data used in this work were acquired at the Helmholtz Zentrum Berlin with the help of N. Kardjilov and A. Hilger as part of the experiment experiment GEO-04-2214. E. Tudisco was supported by the Swedish strategic research programme eSSENCE for this work.

References

[1] E. Tudisco, S.A. Hall, E.M. Charalampidou, N. Kardjilov, A. Hilger, and H. Sone. Full-field measurements of strain localisation in sandstone by neutron tomography and 3D-volumetric digital image correlation. *Physics Procedia*, 69:509–515, 2015.

[2] R. Xia, J. Zhao, and Y. Liu. A robust feature-based registration method of multimodal image using phase congruency and coherent point drift. In *Eighth International Symposium on Multispectral Image Processing and Pattern Recognition*, pages 891903–891903. International Society for Optics and Photonics, 2013.

[3] D.K. Sahu and M.P. Parsai. Different image fusion techniques–a critical review. *International Journal of Modern Engineering Research*, 2(5):4298–4301, 2012.

[4] C.S. Xydeas and V. Petrovic. Objective image fusion performance measure. *Electronics letters*, 36(4):308–309, 2000.

[5] A. Papoulis and S.U. Pillai. *Probability, random variables, and stochastic processes*. Tata McGraw-Hill Education, 2002.

[6] B. Zitova and J. Flusser. Image registration methods: a survey. *Image and vision computing*, 21(11):977–1000, 2003.

[7] J.P.W. Pluim, J.B.A. Maintz, and M.A. Viergever. Image registration by maximization of combined mutual information and gradient information. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 452–461. Springer, 2000.

[8] P. Viola and W.M. Wells III. Alignment by maximization of mutual information. *International journal of computer vision*, 24(2):137–154, 1997.

[9] T. Gaens, F. Maes, D. Vandermeulen, and P. Suetens. Non-rigid multimodal image registration using mutual information. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 1099–1106. Springer, 1998.

[10] A. Roche, G. Malandain, X. Pennec, and N. Ayache. The correlation ratio as a new similarity measure for multimodal image registration. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 1115–1124. Springer, 1998.

[11] E. D'agostino, F. Maes, D. Vandermeulen, and P. Suetens. A viscous fluid model for multimodal non-rigid image registration using mutual information. *Medical image analysis*, 7(4):565–575, 2003.

[12] C. Chefd'Hotel, G. Hermosillo, and O. Faugeras. Flows of diffeomorphisms for multimodal image registration. In *Biomedical Imaging, 2002. Proceedings. 2002 IEEE International Symposium on*, pages 753–756. IEEE, 2002.

[13] H. Schreier, J.J. Orteu, and M.A. Sutton. *Image Correlation for Shape, Motion and Deformation Measurements: Basic Concepts, Theory and Applications*. Boston, MA: Springer-Verlag US, 2009.

[14] P.K. Rastogi and E. Hack. *Optical methods for solid mechanics: a full-field approach*. John Wiley & Sons, 2012.

[15] A. Guery, F. Latourte, F. Hild, and S. Roux. Characterization of SEM speckle pattern marking and imaging distortion by digital image correlation. *Measurement Science and Technology*, 25(1):015401, 2013.

[16] A. Kaestner, M. Morgano, J. Hovind, and E. Lehmann. Bimodal imaging using neutrons and X-rays. 2015.

[17] A. Kaestner, D. Mannes, J. Hovind, P. Boillat, and E. Lehmann. Combined neutron and X-ray imaging on different length scales. 2016.

[18] J.J. Orteu. 3-D computer vision in experimental mechanics. *Optics and Lasers in Engineering*, 47(3):282–291, 2009.

[19] A. Charbal, J.E. Dufour, F. Hild, M. Poncelet, L. Vincent, and S. Roux. Hybrid stereocorrelation using infrared and visible light cameras. *Experimental Mechanics*, 56(5):845–860, 2016.

[20] R.A. Brooks and G. Di Chiro. Beam hardening in x-ray reconstructive tomography. *Physics in medicine and biology*, 21(3):390, 1976.

[21] P. Helnwein. Some remarks on the compressed matrix representation of symmetric second-order and fourth-order tensors. *Computer methods in applied mechanics and engineering*, 190(22):2753–2770, 2001.

[22] B. Wagne, S. Roux, and F. Hild. Spectral approach to displacement evaluation from image analysis.

The European Physical Journal Applied Physics, 17(3):247–252, 2002.

[23] F. Hild, B. Raka, M. Baudequin, S. Roux, and F. Cantelaube. Multiscale displacement field measurements of compressed mineral-wool samples by digital image correlation. *Applied optics*, 41(32):6815–6828, 2002.

[24] S. Roux, F. Hild, P. Viot, and D. Bernard. Three-dimensional image correlation from X-ray computed tomography of solid foam. *Composites Part A: Applied science and manufacturing*, 39(8):1253–1265, 2008.

[25] P. Burt and E. Adelson. The Laplacian pyramid as a compact image code. *IEEE Transactions on communications*, 31(4):532–540, 1983.

[26] E. Klein, P. Baud, T. Reuschlé, and T.F. Wong. Mechanical behaviour and failure mode of Bentheim sandstone under triaxial compression. *Physics and Chemistry of the Earth, Part A: Solid Earth and Geodesy*, 26(1-2):21–25, 2001.

[27] A. Hilger, N. Kardjilov, M. Strobl, W. Treimer, and J. Banhart. The new cold neutron radiography and tomography instrument CONRAD at HMI Berlin. *Physica B: Condensed Matter*, 385:1213–1215, 2006.

[28] A. Tengattini, D. Atkins, B. Giroud, E. Andò, J. Beaucour, and G. Viggiani. NeXT-Grenoble, a novel facility for Neutron and X-ray Tomography in Grenoble. ICTMS2017, 2017.

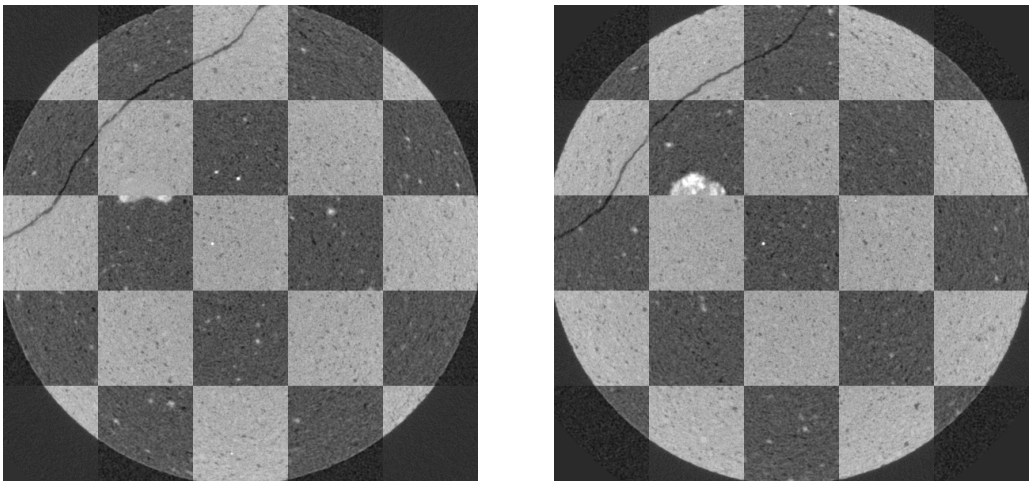

(a) (b)

Figure 7. Complementary composite images formed by a checkerboard assembly of squares extracted from the two modalities after registration. The absence of discontinuity can be checked along all boundaries