



LJMU Research Online

Baker, JL and Einav, I

Deep velocimetry: Extracting full velocity distributions from projected images of flowing media

<http://researchonline.ljmu.ac.uk/id/eprint/17852/>

Article

Citation (please note it is advisable to refer to the publisher's version if you intend to cite from this work)

Baker, JL and Einav, I (2021) Deep velocimetry: Extracting full velocity distributions from projected images of flowing media. Experiments in Fluids, 62 (5). ISSN 0723-4864

LJMU has developed **LJMU Research Online** for users to access the research output of the University more effectively. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LJMU Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain.

The version presented here may differ from the published version or from the version of the record. Please see the repository URL above for details on accessing the published version and note that access may require a subscription.

For more information please contact researchonline@ljmu.ac.uk

<http://researchonline.ljmu.ac.uk/>



Deep velocimetry: Extracting full velocity distributions from projected images of flowing media

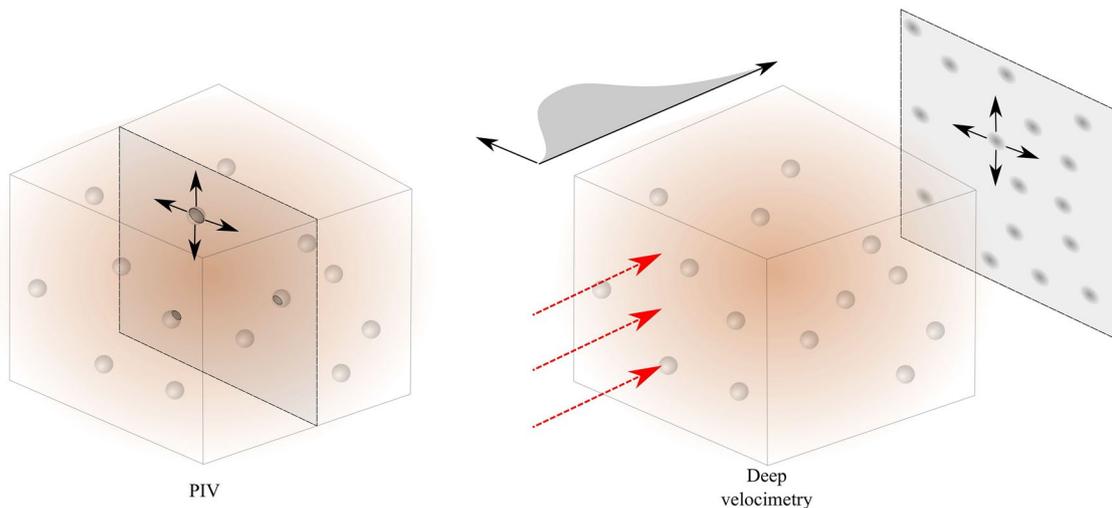
James Lindsay Baker¹ · Itai Einav¹

Received: 8 October 2020 / Revised: 7 March 2021 / Accepted: 5 April 2021 / Published online: 22 April 2021
© The Author(s) 2021

Abstract

Particle image velocimetry (PIV) is a powerful image correlation method for measuring bulk velocity fields of flowing media. It typically uses optical images, representing quasi-two-dimensional experimental slices, to measure a single velocity value at each in-plane position. However, projection-based imaging methods, such as x-ray radiography or shadowgraph imaging, encode additional out-of-plane information that regular PIV is unable to capture. Here, we introduce a new image analysis method, named deep velocimetry, that goes beyond established PIV methods and is capable of extracting full velocity distributions from projected images. The method involves solving a deconvolution inverse problem to recover the distribution at each in-plane position, and is validated using artificial data as well as controlled laboratory x-ray experiments. The additional velocity information delivered by deep velocimetry could provide new insight into a range of fluid and granular flows where out-of-plane variation is significant.

Graphic abstract



1 Introduction

Since its initial development in the 1980s, particle image velocimetry (PIV) has proved to be an invaluable flow measurement tool, and is now arguably the dominant velocimetry technique in experimental fluid mechanics (Westerweel et al. 2013). This is due to its ability to unobtrusively measure instantaneous velocity fields at high spatial and temporal

✉ James Lindsay Baker
james.baker@sydney.edu.au

¹ School of Civil Engineering, The University of Sydney, Sydney, NSW 2006, Australia

resolutions, which has been facilitated by rapid advances in sensor hardware and computer processing power. In the most common application of PIV, a transparent flow is seeded with opaque tracer particles and a laser sheet is used to illuminate an internal slice of a given experiment (Fig. 1a). A high-speed camera then captures two images in quick succession, each of which are split into a series of discrete interrogation windows. By computing the peak of the cross-correlation function in these windows, it is possible to deduce the most likely in-plane particle displacement between the two images, hence providing two-dimensional velocity fields at each time and spatial location.

The accuracy of this classical PIV method has been thoroughly investigated, resulting in a series of guidelines about the interrogation window size, in-plane velocity gradients, out-of-plane motion and seeding particle density for reliable fluid flow measurements (Keane and Adrian 1992). Furthermore, the same principles of PIV can also be applied to optical imaging of dense granular flows (Lueptow et al. 2000), which have the advantage that the grains already provide sufficient contrast without requiring additional flow seeding. However, unlike regular fluids, PIV of granular materials is typically restricted to flow boundaries due to the opaque nature of grains, and therefore not necessarily representative of the bulk. One method to overcome this difficulty is to use refractive index matched scanning (e.g. Dijkstra et al. 2012) to illuminate internal slices, but this involves the introduction of interstitial fluid that changes the flow dynamics.

For flows of both fluids and grains, PIV measurements of velocities at different spatial points can be used to derive

additional quantities of interest such as strain-rate and vorticity fields. Similarly, a sufficient quantity of transient measurements can be used to compute mean and fluctuating velocity components, which provide in-plane Reynolds stresses (in fluids) and granular temperature measurements (in particulate systems). These statistics are usually obtained by computing the instantaneous velocity at each time step and assimilating the ensemble properties of such measurements (e.g. Sarno et al. 2018). However, if only the steady-state fields are required then more accurate results can be obtained by averaging all the cross-correlation functions before finding the single peak, representing the mean velocity in each window over time (Delnoij et al. 1999; Meinhart et al. 2000). This reduces the signal-to-noise ratio, and allows meaningful results to be obtained even when windows of a single pixel are used.

In fact, this average cross-correlation function has been shown to actually represent the convolution of the velocity probability distribution function (taken over all times and in-plane positions within the window) with the particle image function (Westerweel 2008). Consequently, deconvolution can be used to recover the full velocity distribution, which has been demonstrated to good effect when computing in-plane Reynolds stresses in turbulent flows (Scharnowski et al. 2012).

In all of the above, each interrogation window gives rise to a single spatial measurement, and the effect of spatial gradients is usually minimised or explicitly quantified and corrected for in the final measurements. However, it is possible to recover useful spatial information within an individual

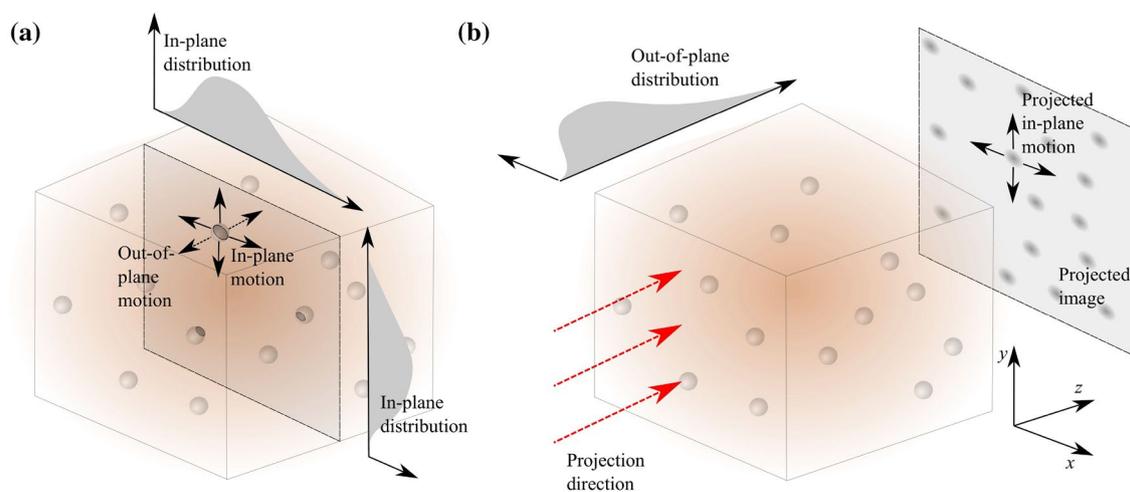


Fig. 1 Schematic of imaging setups for **a** regular particle image velocimetry (PIV) and **b** the projection-based imaging employed in this paper. In **a** a single two-dimensional slice of a three-dimensional volume is imaged. Material motion in the two in-plane directions is captured, with out-of-plane motion taking material in and out of view. Variation along the two in-plane directions gives rise to in-plane dis-

tributions. On the other hand, the projection imaging system **b** incorporates material from the whole out-of-plane path of the projection, although the final image only captures motion in the in-plane directions. In this case, variation along the out-of-plane direction gives rise to out-of-plane distributions

cell. For example, by assuming Gaussian velocity distributions (in space) and Gaussian auto- and cross-correlation peaks, simultaneous measurements of both the mean velocity and granular temperature, or velocity fluctuations, can be made in each window (Qiao et al. 2007).

All of these classical PIV approaches use optical images, and the resulting measurements are confined to in-plane velocity components on a two-dimensional experimental section (Fig. 1a). For transparent fluids, there are a number of ways of extending this planar PIV to obtain additional information. For example, by using two cameras focused on the same plane, stereoscopic PIV (Prasad 2000) allows all three components of the velocity to be computed on that plane, and dual-plane stereoscopic PIV (Kähler and Kompenhans 2000), where two additional cameras are focused on an adjacent plane, gives all components of the velocity and velocity gradient tensor on a single plane. Other more sophisticated techniques include scanning planar PIV (Brücker 1997), tomographic PIV (Elsinga et al. 2006) and holographic PIV (Katz and Sheng 2010), which use various multi-plane approaches to first recover volumetric data that is then correlated to deduce fully three-dimensional displacements. Alternatively, volumetric particle tracking velocimetry (e.g. Bhattacharya and Vlachos 2020) uses multiple cameras to triangulate the three-dimensional position of individual particles at each time step. Many of these techniques, however, can be expensive to implement and difficult to achieve highly accurate measurements with. Furthermore, in certain situations, especially in-situ field measurements, it is simply not possible to position lasers so that they illuminate arbitrary planes, due to obstructions and other practical constraints.

A major limitation of all of the above advanced PIV techniques is that they require optical imaging of internal regions of the flow. This is possible for transparent fluids, but not for opaque granular materials. For this reason, a number of unobtrusive imaging methods have been applied to flowing granular materials, for example magnetic resonance imaging (MRI) (Stannarius 2017) or positron emission particle tracking (PEPT) (Parker 2017). X-ray computed tomography (CT) (Hall et al. 2010) is particularly promising, because it is applicable to general materials and offers high spatial resolution. However, to produce detailed 3D density maps, x-ray CT requires large numbers of projections, or radiographs, from different scanning directions, hence severely limiting its temporal resolution.

These temporal resolution problems can be avoided by working directly with a single set of x-ray projections, since it is possible to continuously acquire images from a fixed direction as material flows or deforms, which is called dynamic x-ray radiography. By applying similar principles to planar PIV and correlating subsequent images, the cross-correlation peak generally gives the most likely in-plane

velocity in each interrogation window. Unlike planar imaging, the integrated nature of x-ray radiography means that these velocities are actually representative of the whole depth of the sample, not just at the boundaries (Fig. 1b), and thus can be considered beam-averaged quantities. This approach has been used to successfully measure such velocities in both fluid (Kim and Lee 2006) and granular systems (Guillard et al. 2017).

Many flows are not satisfactorily described using single beam-averaged velocities in each interrogation window, especially fully three-dimensional systems with strong variation in the out-of-plane direction. In reality, there is actually a distribution of velocities through the path of the x-ray beam. While each 2D projection does not provide any specific depth information, it does encode additional information about this velocity distribution. Specifically, analogously to the previously mentioned convolution for planar imaging (Westerweel 2008), the cross-correlation function of two successive x-ray radiographs is the convolution of the auto-correlation of the first image with the velocity distribution. This relationship has been exploited to recover crucial additional velocity information in fluids (Dubsky et al. 2010) and granular materials (Baker et al. 2018). Combining with multi-directional x-ray imaging, both of these approaches were able to measure three-dimensional velocity fields at every point in space, a significant feat for optically opaque experiments.

This paper further builds on the concept of out-of-plane velocity distributions for x-ray radiography. We present a simplified convolution equation relating the projected image intensities to the velocity distribution, without the need to first compute any correlation functions. This relationship can be efficiently inverted to recover the full velocity distribution in each interrogation window, a process we call deep velocimetry due to the additional depth-information provided over classical planar PIV. We believe that this extra capability is a simple yet powerful complementary method to PIV and its various extensions. It is motivated by x-ray projections, but could equally be applied to other projection-based imaging methods, such as shadowgraph or schlieren imaging (Luthman et al. 2019; Ozawa et al. 2020), which are often cheaper and more robust than more complicated multi-dimensional methods. We therefore believe that deep velocimetry has potential applications to a range of granular and fluid flow systems. Indeed, the generality of the method makes it suitable for any heterogeneous media with sufficient macroscopic texture, for example foams and various types of soils.

The rest of this paper is organised as follows: in Sect. 2 we introduce the theory behind the new deep velocimetry method. Section 3 then demonstrates its effectiveness and robustness compared to existing methods on a simple one-dimensional test case, before Sect. 4 applies it to a more realistic example

representing x-ray imaging of granular materials. Experimental validation in the form of x-rays of controlled motion of granular materials is provided in Sect. 5. Section 6 discusses potential applications to other imaging systems, and Sect. 7 draws final conclusions.

2 Theory of deep velocimetry

In this section we introduce the theory behind the new method for an idealised projection-based imaging system, which is assumed to directly measure the integrated density of material along the projection direction. To transfer these results to real imaging systems requires knowledge of the relationship between this integrated value and the recorded projections, which will depend on the specific system and will be discussed later. The theory is restricted to a single pair of projected images—an *initial* projection and a *moved* projection a finite time later—and is used to recover the displacements between these images. In practice, these displacements will usually be converted to velocities using a known time difference, and a sequence of many image pairs is typically used to compute transient measurements or accurate mean quantities.

For a single pair, let $\rho_1(x, y, z)$ denote the initial density field of the flowing material, where (x, y, z) represents a standard Cartesian coordinate system in three-dimensional space. As indicated on Fig. 1b), we assume that the projections correspond to parallel rays in the z direction, with x and y representing the two directions perpendicular to the projections. The initial integrated density, or depth of material D_1 , is then defined as

$$D_1(\mathbf{x}) = \int_{z_{min}}^{z_{max}} \rho_1(x, y, z) dz, \tag{1}$$

where z_{min} and z_{max} represent the limits of the projection and $\mathbf{x} = (x, y)$ denotes the two projected in-plane directions. Similarly, we can define the moved density field $\rho_2(x, y, z)$ and corresponding depth D_2 as

$$D_2(\mathbf{x}) = \int_{z_{min}}^{z_{max}} \rho_2(x, y, z) dz. \tag{2}$$

Between these two instances, it is assumed that the material has moved in the two in-plane directions (x, y) according to the displacement field $\mathbf{u} = (u, v)$ which will, in general, depend on all spatial coordinates (x, y, z) . There may also be motion in the out-of-plane (z) direction, but this is not important as parallel projections are unable to depict such motion, providing all material remains confined between the limits z_{min} and z_{max} . At each in-plane position \mathbf{x} , the displacement field \mathbf{u} gives rise to a probability density function (PDF) $f_x(\mathbf{v})$, which reflects the fact that it is not possible

to determine the relative out-of-plane position of material when using projection imaging. This PDF represents the likelihood of finding a particular displacement \mathbf{v} at a pre-determined in-plane position \mathbf{x} yet a randomly chosen out-of-plane z position, and is given by the explicit expression

$$f_x(\mathbf{v}) = \frac{\int_{z_{min}}^{z_{max}} \rho_1(x, y, z) \delta(\mathbf{u}(x, y, z) - \mathbf{v}) dz}{\int_{z_{min}}^{z_{max}} \rho_1(x, y, z) dz}, \tag{3}$$

where $\delta(\mathbf{s})$ is the extension of the Dirac delta function to two-dimensional domains. Eq. (3) is nonnegative since the initial density field ρ_1 cannot be negative, and it is also straightforward to see that

$$\int_{u_{min}}^{u_{max}} \int_{v_{min}}^{v_{max}} f_x(\mathbf{v}) d\mathbf{v} = 1, \tag{4}$$

where u_{min} , u_{max} , v_{min} and v_{max} are the minimum and maximum displacements in the two in-plane \mathbf{x} directions, taken over all out-of-plane z locations. Thus, (3) does indeed define a probability density function related to the barycentric (density-weighted) velocity profile. Noting that the denominator of (3) is equal to $D_1(\mathbf{x})$, it follows that

$$D_1(\mathbf{x})f_x(\mathbf{v}) = \int_{z_{min}}^{z_{max}} \rho_1(x, y, z) \delta(\mathbf{u}(x, y, z) - \mathbf{v}) dz, \tag{5}$$

which represents the integrated density of material that moves from initial position \mathbf{x} by displacement \mathbf{v} . This can be related to the moved projection D_2 as follows. At a given in-plane position \mathbf{x} of the moved image, the integrated density of material that has moved by displacement \mathbf{v} is simply the amount of material that has moved from position $\mathbf{x} - \mathbf{v}$ by \mathbf{v} , i.e.

$$D_1(\mathbf{x} - \mathbf{v})f_{\mathbf{x}-\mathbf{v}}(\mathbf{v}). \tag{6}$$

The moved image D_2 is then given by the integral of all such displacements, namely

$$D_2(\mathbf{x}) = \int_{u_{min}}^{u_{max}} \int_{v_{min}}^{v_{max}} D_1(\mathbf{x} - \mathbf{v})f_{\mathbf{x}-\mathbf{v}}(\mathbf{v}) d\mathbf{v}. \tag{7}$$

Hence, there is a clear relationship between the initial and moved projections and the displacement PDF at each in-plane position. In general, it is not possible to use expression (7) to recover every PDF, which would require using a function of only one variable (\mathbf{x}) to solve for a function of two variables (\mathbf{x} and \mathbf{v}). However, much progress can be made by making one key additional observation. This is achieved by considering finite two-dimensional interrogation windows of the projected images. In an analogous manner to the representative volume elements typically employed in

continuum mechanics, it is assumed that material properties do not vary significantly within this interrogation window and can be replaced by single quantities that are representative of the whole window. For the case of projected images, the interrogation window W needs to be chosen so that the velocity profile does not have any strong in-plane spatial gradients within the window, and also to ensure that the density profile is approximately constant at different \mathbf{x} positions within the bounds that define the window. With a suitable choice of interrogation window, these facts ensure that the resulting PDFs $f_x(\mathbf{v})$ can be replaced by a single representative PDF $f_W(\mathbf{v})$ in each window W , which is independent of the local in-plane position within the window. To see this, we can replace the fully spatially dependent densities and velocities in the PDF definition (3) by values that depend on z only. The resulting depth integrals are thus independent of in-plane position within each window. Note that we do not require homogenous density or velocity profiles in the out-of-plane direction. Indeed, it is such non-uniform behaviour that leads to more complicated and interesting distributions.

For the remainder of this paper we will only consider a single interrogation window, and therefore we drop the W subscript and assume

$$f_x(\mathbf{v}) \approx f(\mathbf{v}), \tag{8}$$

where $f(\mathbf{v})$ is the representative, or mean, displacement PDF of the whole in-plane domain. With this approximation, equation (7) simplifies to

$$D_2(\mathbf{x}) \approx \int_{u_{min}}^{u_{max}} \int_{v_{min}}^{v_{max}} D_1(\mathbf{x} - \mathbf{v})f(\mathbf{v}) \, d\mathbf{v}, \tag{9}$$

which can be written as a convolution equation

$$D_2 \approx D_1 * f, \tag{10}$$

where ‘*’ is the convolution operator. This relates the mean PDF f to projected images D_1 and D_2 . Since these images are both given, solving the deconvolution inverse problem is a relatively straightforward way to calculate the displacement PDF.

2.1 Relationship to correlation functions

Note that this deconvolution approach is slightly different to that taken by Baker et al. (2018), who first calculated the auto- and cross-correlation functions of their projected images before computing a deconvolution to deduce the displacement PDF. We will show in the next two sections that working directly with the images provides accurate results without requiring this additional step but, for completeness, we now briefly explain why such a correlation-based approach also works from a theoretical perspective. The

auto-correlation function A and cross-correlation function C are defined in terms of the projected depths D_1 and D_2 as

$$A = D_1 \star D_1, \quad C = D_1 \star D_2, \tag{11}$$

where

$$A(\mathbf{s}) = \int \int D_1(\mathbf{x})D_1(\mathbf{s} + \mathbf{x}) \, d\mathbf{x}, \tag{12}$$

$$C(\mathbf{s}) = \int \int D_1(\mathbf{x})D_2(\mathbf{s} + \mathbf{x}) \, d\mathbf{x}, \tag{13}$$

and the integrals are taken over the appropriate representative interrogation window. Next, if we assume that D_1 , D_2 and f are related by the convolution Eq. (10) then, using the basic properties of correlations and convolutions, it is straightforward to see that

$$C = D_1 \star D_2 = D_1 \star (D_1 * f) = (D_1 \star D_1) * f = A * f. \tag{14}$$

Hence, A , C and f are also related by an identical convolution equation $C = A * f$, meaning the auto- and cross-correlation functions can alternatively be deconvolved to recover the velocity distribution.

3 Example 1: One-dimensional test case

3.1 Setup

The capability of this new deep velocimetry method is first demonstrated on a simple system that consists of two-dimensional density fields $\rho(x, z)$, which are projected onto one-dimensional images $D(x)$. The setup is shown on Fig. 2, where it is assumed that the initial density field $\rho_1(x_i, z_j)$ is defined at N_x in-plane positions x_i and N_z out-of-plane positions z_j , which are both equally spaced between zero and one. The density field is equal to either zero or one at each grid point (Fig. 2a), representing a system consisting of uniform solid material or void space. The exact configuration is determined by choosing the nonzero positions from a uniform distribution and specifying a mean solid volume fraction

$$\bar{\phi} = \frac{1}{N_x N_z} \sum_{i=1}^{N_x} \sum_{j=1}^{N_z} \rho_1(x_i, z_j). \tag{15}$$

This initial distribution is then subjected to a one-dimensional displacement field $u(z)$ in the x direction, which is chosen to be the quadratic profile

$$u(z) = 6\bar{u}z(1 - z), \tag{16}$$

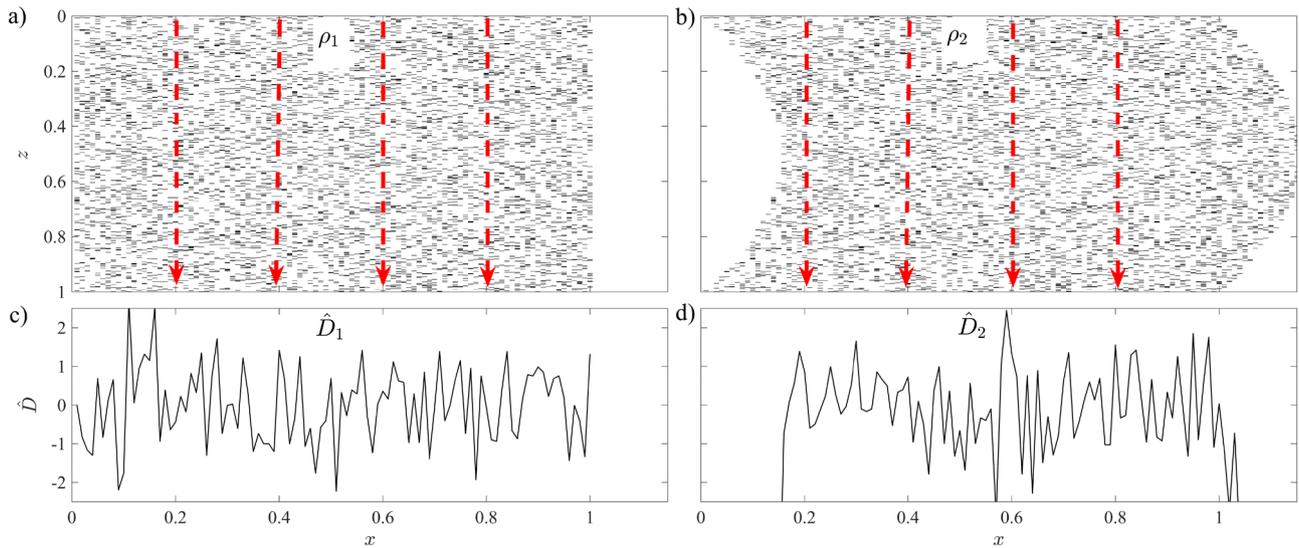


Fig. 2 Example density and projected images for the one-dimensional test case of Sect. 3, showing **a** the initial density field $\rho_1(x, z)$, **b** the moved density field $\rho_2(x, z)$, and **c**, **d** the corresponding normalised initial and moved projected depths $\hat{D}_1(x)$ and $\hat{D}_2(x)$ (for visualisation, dimensional values D_1 and D_2 are scaled to have mean zero and

standard deviation one). In **a** and **b** black regions denote a density of one, and white corresponds to a density of zero. Red arrows denote expected projection direction. Parameters chosen are $N_x = 100$, $N_z = 10000$, $\bar{\phi} = 0.1$ and $\bar{u} = 10$ px

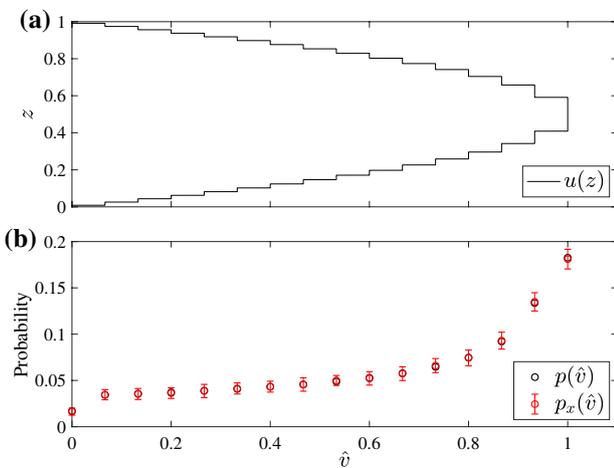


Fig. 3 Plots of **a** the imposed displacement profile $v = u(z)$ from (16) and **b** the corresponding probability mass functions (PMFs). Black circles in **b** represent the theoretical PMF $p(v)$ from equation (18), and red error bars are the mean and standard deviation of the actual realisations p_{x_i} over the N_x different x positions. Displacements have been normalised to have maximum one, with other parameters the same as Fig. 2

where \bar{u} represents the mean displacement magnitude ($\int_{z_{min}}^{z_{max}} u(z) dz = \bar{u}$). Displacements are rounded to whole pixels (Fig. 3a), and the density field is then moved according to these values to produce the moved density field ρ_2 (Fig. 2b), which also consists exclusively of ones and zeros.

The projected normalised depths \hat{D}_1 and \hat{D}_2 of these density fields are shown on Figs. 2c and d, respectively, which represent the dimensional quantities D_1 and D_2 scaled to have mean zero and standard deviation one. This scaling is used solely for visualisation in this paper, but can also be useful to remove the effects of spatial gradients and background noise in real imaging systems.

At each in-plane position x_i we can define a velocity distribution in an analogous manner to (4) for discrete positions, namely

$$p_{x_i}(v) = \frac{1}{D_1(x_i)} \sum_{j=1}^{N_z} \rho_1(x_i, z_j) \delta(u(z_j) - v), \tag{17}$$

where $\delta(s) = 1$ when $s = 0$ and $\delta(s) = 0$ otherwise. In this discrete case, Eq. (17) defines probability mass functions (PMFs) that each sum to unity, and Fig. 3b shows plots of these PMFs. The error bars on Fig. 3b show that there is some scatter for different x positions but the PMFs are well approximated by the mean function, taken over all positions x_i .

To further investigate the properties of the velocity distribution (17) it is useful to think in terms of random variables. Specifically, we can consider only the z positions where the density is nonzero, since these are the only locations that contribute to the velocity distribution, and assume that each of these locations are equally probable during the random generation process. If there are N_ρ of these nonzero positions, we can treat the specific location of each as random variables

$Z_1, Z_2, \dots, Z_{N_\rho}$, and introduce the multivariate random variable $Z = (Z_1, Z_2, \dots, Z_{N_\rho})$ to represent the full vector of nonzero positions. At each position, the PMF (17) then corresponds to a particular realisation of the random variable Z . We omit the derivation for brevity, but it can be shown, using the properties of random variables, that the expected value of the PMF is then

$$p(v) = \frac{1}{N_z} \sum_{j=1}^{N_z} \delta(u(z_j) - v), \tag{18}$$

which assumes that every z point contributes equally to the PMF. The standard deviation about this mean can also be computed as

$$\sigma(v) = \sqrt{\frac{1 - \phi}{\phi} \frac{1}{N_z - 1} p(v)(1 - p(v))}, \tag{19}$$

where $\phi = N_\rho / N_z$ is the solid volume fraction at a given x position. Note that $\sigma \rightarrow 0$ as $\phi \rightarrow 1$ or $N_z \rightarrow \infty$, meaning that fluctuations from the mean PMF at each spatial position become insignificant. Hence, the assumption that the PMF is independent of in-plane position becomes increasingly valid as we approach these limits. Conversely, it is expected that the assumption will break down in sparse systems as the volume fraction ϕ approaches zero, or in shallow systems that represents projections over few elements N_z . Both Eqs. (18) and (19) have been verified numerically by computing many randomly generated density configurations at different packing fractions and depths.

3.2 Forward convolution

Having established that the displacement PMF is approximately independent of in-plane position (within a representative interrogation window), we next test the result (10) that the moved projection is the convolution of the initial projection with the PMF. In the discrete, one-dimensional, case this can be written

$$D_2 = D_1 * p, \tag{20}$$

where the discrete convolution is defined as

$$(D_1 * p)(x_i) = \sum_{j=1}^{N_v} D_1(x_i - v_j) p(v_j), \tag{21}$$

for N_v distinct displacements. Figure 4 shows an example comparison between D_2 and $D_1 * p$, with excellent agreement confirming that this is indeed an appropriate approximation for the particular set of parameters chosen. In general, the validity of the approximation will again depend on the specific parameters used. From (19) it follows that ϕ and

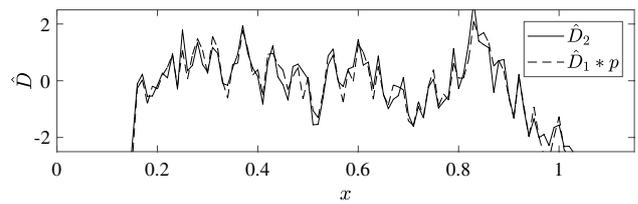


Fig. 4 Plots of the actual (normalised) moved image \hat{D}_2 (solid line) and the predicted forward convolution ($\hat{D}_1 * p$) from Eq. (21) (dashed line). Other parameters are the same as Figs. 2 and 3

N_z again play an important role, although the exact dependence is non-trivial due to the nonlinearity of the convolution operator (21).

3.3 Deconvolution

Finally, we use the two projections D_1 and D_2 to recover the PMF p by solving a deconvolution inverse problem. There are a number of different ways of achieving this, but here we choose to solve a constrained linear least-squares approach by first recasting (21) as the matrix equation

$$Ax = b, \tag{22}$$

where

$$A = \begin{pmatrix} 1 & 1 & \dots & 1 \\ D_1(N_v) & D_1(N_v - 1) & \dots & D_1(1) \\ D_1(N_v + 1) & D_1(N_v) & \dots & D_1(2) \\ \vdots & \ddots & \ddots & \vdots \\ D_1(N_x) & D_1(N_x - 1) & \dots & D_1(N_x - N_v + 1) \end{pmatrix}, \tag{23}$$

$$x = (p(1), p(2), \dots, p(N_v))^T, \tag{24}$$

$$b = (1, D_2(1), D_2(2), \dots, D_2(N_x - N_v + 1))^T. \tag{25}$$

Note that the top row of the matrix A has been introduced to ensure that the probabilities sum to one. Eq. (22) is then solved using the *lsqlin* function in Matlab, subject to the additional nonnegativity constraint of probability mass functions, $x(i) \geq 0$ for all i . This is slightly different to the deconvolution approach adopted by Baker et al. (2018), who solved a constrained nonlinear optimisation problem with an additional Tikhonov regularisation to ensure that the distribution did not vary in an unphysical manner. Section 3.4 compares the results of these two methods in more detail. Note that a third possible deconvolution method involves solving the unconstrained linear system (22)–(25) directly using, for example, Gaussian elimination. This is a constructive method that requires the least computational time and was found to sometimes give very accurate results. However, since it does not explicitly constrain the system to

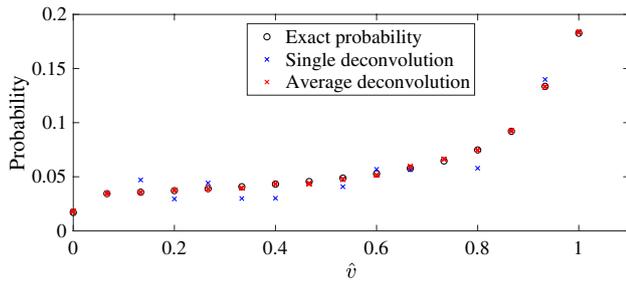


Fig. 5 Plots of the exact underlying PMF (black circles), reconstructed PMF from a single deconvolution (blue crosses) and average reconstructed PMF by taking the mean of 100 deconvolutions (red crosses). All parameters are same as previous figure

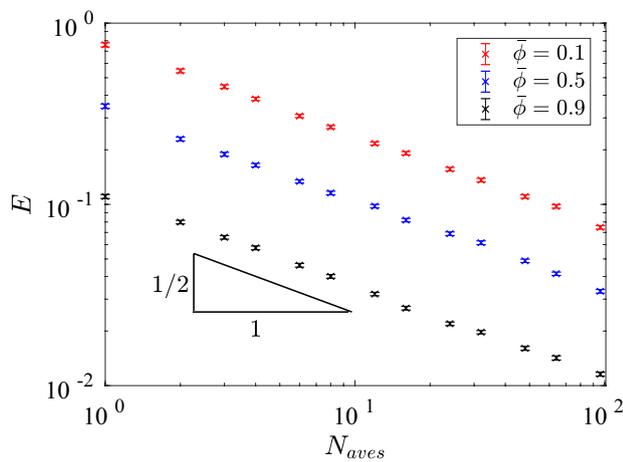


Fig. 6 Plots of the error (26) as a function of number of deconvolution averages for three different mean solid fractions $\bar{\phi}$. Error bars are calculated from computing 100 different results for each point. All other parameters are same as previous figure

nonnegative values, it sometimes produces unphysical PMFs that are less than zero, and we therefore do not employ this direct approach.

Figure 5 shows an example of the reconstructed PMF, which is in good agreement with the exact underlying distribution $p(v)$. This demonstrates the effectiveness of the new deep velocimetry method when deconvolving a single pair or projected images. However, in practise the reliability of individual measurements can be improved by taking the average of multiple instances. To this end, Fig. 5 also shows the resulting PMF after computing 100 distinct deconvolutions, with differing density configurations, and then taking the mean of these PMFs. This shows a significant improvement, with the final PMF being almost indistinguishable from the exact underlying distribution.

To investigate the effect of averaging further, Fig. 6 shows the final deconvolution error

$$E = \sum_{j=1}^{N_v} |p(v_j) - \tilde{p}(v_j)|, \tag{26}$$

as a function of the number of averages N_{aves} , where $\tilde{p}(v)$ represents the approximation to the true $p(v)$ by averaging the deconvolutions. We see that this error consistently decreases with increasing degree of averaging. In fact, it appears that $E \propto N_{aves}^{-1/2}$ for the range of parameters explored here, although the origin and full extent of this scaling law remains to be investigated. It is also clear that higher solid volume fractions $\bar{\phi}$ give more accurate results, which is consistent with the earlier observation (19) that the initial PMF is closer to the mean value at each spatial position in these denser packed regimes. In all cases, the method rapidly produces satisfactory errors of less than 10%. This one-dimensional test case therefore demonstrates the potential of deep velocimetry as a highly accurate method of measuring full out-of-plane velocity distributions.

3.4 Comparison to correlation-based methods

In this section we compare the results of deep velocimetry to previous correlation-based methods, such as those employed by Dubsy et al. (2010) and Baker et al. (2018), to assess the relative merits of this new technique.

The correlation-based approaches involve first computing the auto-correlation function of an initial image, as well as the cross-correlation function of this image with the moved image. As described in Sect. 2.1, these correlation functions are related to each other and the displacement probability density function f by a similar convolution Eq. (14). We follow the procedure described in Baker et al. (2018) to solve this deconvolution problem and extract f for the same one-dimensional test case. When using correlation functions, there are actually two different means of averaging over a number of distinct results. The first is to follow the same averaging procedure as adopted above for deep velocimetry and compute N_{aves} distinct deconvolutions. These PMFs are then averaged to give the final result. This method will be referred to as ‘post-averaging’, because the averaging only takes place at the final stage. Alternatively, it is possible to ‘pre-average’ by first computing a number of auto- and cross-correlation functions and averaging these functions. The deconvolution is then only carried out once on these average correlation functions. The rationale behind such an approach is that deconvolution of the average functions is less sensitive to noise than deconvolving a single noisy function pair. Note this pre-averaging strategy is not a viable option when deconvolving the images directly, because the pre-averaging will smooth out any heterogeneous texture in the images, making deconvolution impossible. The correlation functions, on the other hand, do not tend towards

uniform states upon averaging. Previous authors have employed a combination of both pre- and post-averaging for their correlation-based methods, but in this paper we will focus on pure implementations of both approaches since the optimal averaging combination remains an open question.

Figure 7a shows the deconvolution errors E as a function of the degree of averaging N_{aves} for the three different deconvolution methods, at a fixed solid fraction $\bar{\phi} = 0.5$. They all show a similar trend of decreasing error with N_{aves} , with the same approximate scaling $E \propto N_{aves}^{-1/2}$ for both the direct image deconvolution and correlation-based methods. While all methods provide satisfactory final results, especially when using at least $N_{aves} = 20$ deconvolutions, there is a clear difference between the direct deconvolution and correlation-based approaches. The new method performs consistently better than both of the previous techniques, regardless of the total amount of averaging. After 100 averages, the direct deconvolution approach produces an average residual error of 2.2%, whereas the equivalent post-averaging correlation method is significantly higher at 4.6%. Interestingly, the pre-averaging correlation method is the least accurate for the parameters considered, giving a mean error of 5.6%

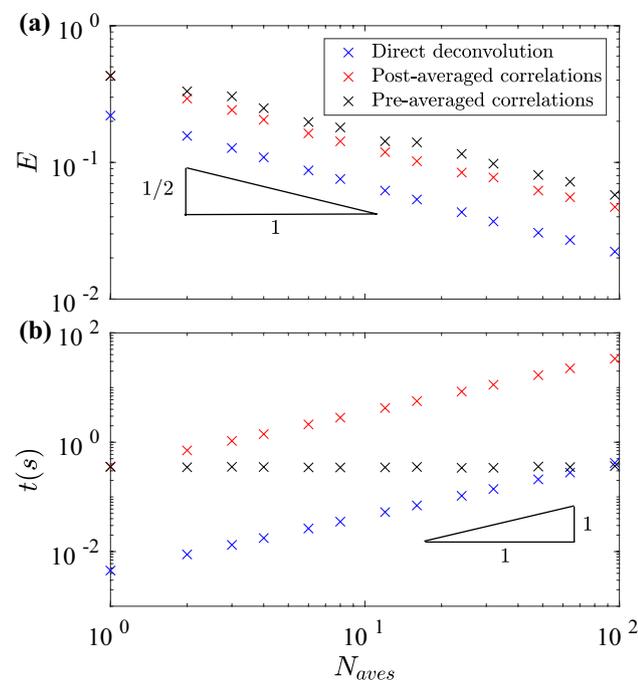


Fig. 7 Plots of **a** the error (26) and **b** the total computational time as a function of number of deconvolution averages. Three different reconstruction methods are employed: direct image deconvolution as described in this paper (blue), deconvolution of auto- and cross-correlation functions with post-averaging of the results (red), and a similar correlation-based method but first pre-averaging the correlation functions (black). Values are calculated from mean of 100 different computations for each point. Mean solid fraction $\bar{\phi} = 0.5$, and all other parameters are same as previous figures

after 100 averages. In general, the difference between the two correlation methods is less prominent than that between the direct deconvolution and the post-averaging correlation approach, although there remains a large region of parameter space yet to be fully explored.

Since reconstruction methods typically involve a trade-off between accuracy and computational time, we also record the CPU time associated with each of the methods. These are plotted on Fig. 7b, and we find that both the direct deconvolution and the post-averaging correlation method display an approximately linear relationship between the amount of averaging and computational time. The direct deconvolution method, however, is significantly faster to implement, requiring approximately 80 times less computational time. Computing the auto- and cross-correlation functions does take additional time, but this is relatively fast. The primary source of the computational time differences is the difference between computing the constrained linear problem considered here compared to the constrained nonlinear optimisation problem of Baker et al. (2018). The latter is much slower and forms the majority of the total computational time for the pre-averaging correlation-based method. The post-averaging method, on the other hand, only requires one deconvolution and is therefore considerably faster than the equivalent post-averaging approach. The total computational time for the pre-averaging method still increases with the amount of averaging, but this is much slower as it is only due to the fast computation of the correlation functions. For low numbers of averages, the pre-averaging approach is slower than direct deconvolution due to the large overhead of a single nonlinear deconvolution. However, as the amount of averaging increases the gap quickly closes. For the parameters used here, by $N_{aves} = 100$ the pre-averaged correlation method takes less time on average than the direct deconvolution, and this will become more exaggerated as the number of averages continues to grow.

In summary, the new direct deconvolution approach appears to perform favourably compared to the previous correlation-based approaches, both in terms of accuracy and total computational time. Depending on the system and resources available, total computational time may not be a limiting factor and hence less relevant when comparing methods, or the required degree of averaging may be such that the pre-averaging correlation approach actually requires less time. Note that this amount of averaging will set the transient timescales that can be measured. In the example above, the underlying motion was assumed identical for each specific realisation, and the final results can therefore be interpreted as steady-state velocity measurements. For flows that evolve over time, on the other hand, the total averaging window needs to be suitably short to capture transient dynamics. This could also affect which of the methods is the most suitable for a given application. Finally, it may also be

possible to improve the speed and accuracy of the correlation-based methods by formulating as constrained linear problems similar to the direct deconvolution approach used here. However, the purpose of this section is to compare our new method to existing reconstruction tools, not to improve such tools. This comparison is very agreeable, suggesting that deep velocimetry offers promising capabilities.

3.5 Effect of noise

The previous results were all constructed using exact image values to test the idealised theory. In reality, this will not be the case. Images will most likely include a certain degree of experimental noise, and deconvolution is known to be sensitive to such noise. To test how deep velocimetry is affected by noise, we repeated a similar set of computations with non-exact, or ‘noisy’ simulated images. Specifically, the initial and moved projected depths, D_1 and D_2 , respectively, were generated as before and then distorted with noise to give the new values, D_1^{noise} and D_2^{noise} , where

$$D_i^{noise}(x) = D_i(x) + N_i(x), \tag{27}$$

for $i = 1, 2$. At each spatial position, the random variables $N_i(x)$ are drawn from a normal distribution of mean zero and standard deviation $\sigma \bar{D}_i$, where \bar{D}_i is the mean projected depth over a given image, and the magnitude σ is systematically varied.

Figure 8 shows the resulting deconvolution errors after $N_{aves} = 100$ computations as a function of the noise magnitude σ , for a fixed packing fraction $\bar{\phi} = 0.5$. For completeness, we compute the errors for all of the three methods described in Sect. 3.4. From Fig. 8, we see that the effect

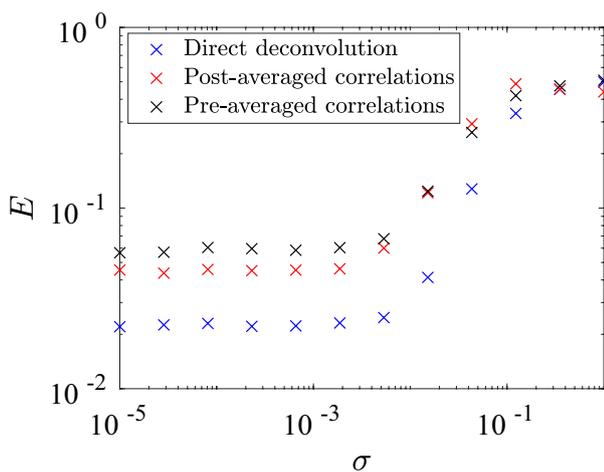


Fig. 8 Plots of the error (26) as a function of the degree σ of added noise for the three different reconstruction methods shown in Fig. 7. Values are calculated from mean of 100 different computations for each point. All other parameters are same as previous figures

of noise is relatively low for small values of σ , with only a gradual decline in accuracy up to $\sigma = 10^{-2}$. Past this point, however, all of the methods start to break down more rapidly, eventually producing large errors for $\sigma \geq 0.1$. As found on Fig. 7a, the direct deconvolution approach again performs equally as well as the two correlation-based approaches. For low amounts of noise, it is significantly more accurate than the previous methods. This gap does begin to close as the amount of noise increases, but the direct deconvolution remains at least as accurate up until the point that none of the methods would be workable in practice. This provides further evidence of the robustness of our new method and comparable performance compared to existing techniques.

Note that the pointwise Gaussian noise introduced for the computations of Fig. 8 is designed to test the robustness of deep velocimetry, and not to precisely match the exact nature of noise encountered for a given experiment. The latter would require more specific knowledge of the particular imaging system and associated noise. For example, systems with long exposure times, such as x-ray radiographs, produce images with significant motion blur. Some correlation-based methods (e.g. Fouras et al. 2007) explicitly account for this form of noise, but we do not attempt to make similar corrections in this preliminary paper. It remains to be investigated whether the direct deconvolution will still outperform the correlation-based approaches in such systems.

4 Example 2: X-rays of granular materials

4.1 X-ray imaging

Next, we move from the idealised one-dimensional setup of Sect. 3 to illustrate how deep velocimetry can be applied to a specific projection-based imaging system, namely dynamic x-ray radiography. Such a system involves an x-ray source that emits approximately parallel rays of a particular initial intensity towards the flowing material. As these rays travel through the material they are attenuated and reduce in intensity. Their final intensity is measured by a detector on the opposite side of the sample, which records projected images, or x-ray radiographs, at specific times.

Unlike the previous example, for such an x-ray system the projected images are not a direct measure of the integrated density field D . Instead, the recorded intensity I at projected in-plane position x typically follows a Beer-Lambert attenuation law

$$I(x) = I_0 \exp\left(-\int_{z_{min}}^{z_{max}} \mu(x, y, z) dz\right), \tag{28}$$

where z_{min} is the out-of-plane location of the x-ray source, z_{max} is the out-of-plane location of the detector, I_0 is the constant initial intensity at the source, and μ is the spatially dependent linear attenuation coefficient. This coefficient depends on the specific material being imaged, as well as other factors such as the x-ray power. For simple binary systems consisting of solid material and interstitial air, the attenuation coefficient of air is several orders of magnitude lower than the solid, and hence can be neglected. If we also assume that all solid regions are made of the same material then we can define a mass coefficient $\mu_m = \mu/\rho$ that is constant in space. The intensity can then be written

$$I(\mathbf{x}) = I_0 \exp\left(-\int_{z_{min}}^{z_{max}} \mu_m \rho(x, y, z) dz\right), \tag{29}$$

which simplifies to

$$I(\mathbf{x}) = I_0 \exp(-\mu_m D(\mathbf{x})), \tag{30}$$

where D is the integrated density field, defined in an analogous manner to (1). Hence, the measured x-ray intensity I is directly related to the integrated density D via the relation

$$I(D) := I_0 \exp(-\mu_m D). \tag{31}$$

Now, recall that the theory of deep velocimetry from Sect. 2 is based around the convolution Eq. (10) for the initial and moved depths, D_1 and D_2 , respectively. Dynamic x-ray imaging does not provide the same depths directly, instead recording the initial and moved intensities, $I_1 := I(D_1)$ and $I_2 := I(D_2)$. A key question therefore arises: can we formulate a similar convolution equation involving these projected intensities?

In general, there is no reason why the relationship $D_2 = D_1 * f$ should imply an equivalent relationship between I_1 , I_2 and f (where f is the probability density function of displacements). However, if we replace the function $I(D)$ from (31) by a simple linear function $I(D) = mD + c$ for constants m and c , then it immediately follows from the distributivity property of the convolution operator that

$$I_2 = I_1 * f. \tag{32}$$

Hence, for linear functions $I(D)$, the problem reduces to solving an identical deconvolution problem to recover the PDF f . Of course, the exact intensity relationship (31) is a nonlinear function, meaning the same logic cannot strictly be applied. However, for many systems the fluctuations in the integrated depth across different spatial positions are typically small compared to the mean value. Thus, providing the attenuation coefficient μ_m is also sufficiently small, Eq. (31) is actually well-approximated by its first-order Taylor expansion about the mean integrated depth. It is therefore

approximately linear, and the convolution equation (32) remains valid.

An alternative method of obtaining an appropriate convolution equation when there are significant depth fluctuations, or large attenuation coefficients, is to take logarithms of the projected images. In this case we work instead with functions $L(D)$, where

$$L(D) := \log(I(D)) = \log I_0 - \mu_m D, \tag{33}$$

is now a linear function of D . Hence, we can use the convolution equation

$$L_2 = L_1 * f, \tag{34}$$

to recover the PDF f , where $L_1 = L(D_1)$ and $L_2 = L(D_2)$. We will show in the next subsection how Eqs. (32) and (34), as well as the original convolution equation (10), can be solved in this manner for a system representing x-rays of flowing granular materials.

4.2 Simulation of granular materials

Figure 9 shows the setup of the simulations, which are conducted in a similar manner to the supplementary information of Baker et al. (2018). A computational box of size $1024 \times 1024 \times 1024$ px is first filled with 10,000 spherical particles of diameter $20 \text{ px} \pm 10\%$. The position and size of these particles is randomly drawn from a uniform distribution across the whole domain. Note that this generation method does potentially allow particles to overlap in an unphysical manner, but this is not expected to significantly influence the final results since, in reality, particles still overlap in the projected images.

These initial particle positions are then used to generate an initial three-dimensional density field ρ_1 , where, as before, ρ_1 is taken to be one at spatial positions inside a given particle and zero otherwise (overlapping grain regions are counted as many times as there are overlapping grains). The integrated initial depth D_1 is then calculated from Eq. (1) using the simple geometry of the spherical grains, and the initial intensity I_1 computed from Eq. (30) with attenuation coefficient $\mu_m = 0.005 \text{ px}^{-1}$ and initial intensity $I_0 = 1$. These projections are shown on Fig. 9b, c, and d shows the log of the intensity (L_1).

Next, the particles are moved according to a prescribed velocity field $u(z)$ in the x direction, where u is now taken to be linear in the out-of-plane direction:

$$u(z) = 2\bar{u}z, \tag{35}$$

with $\bar{u} = 10$ px again representing the mean displacement magnitude. Periodic boundary conditions are prescribed in the x direction, ensuring that particles that leave the domain re-enter at the other boundary. After movement, the new

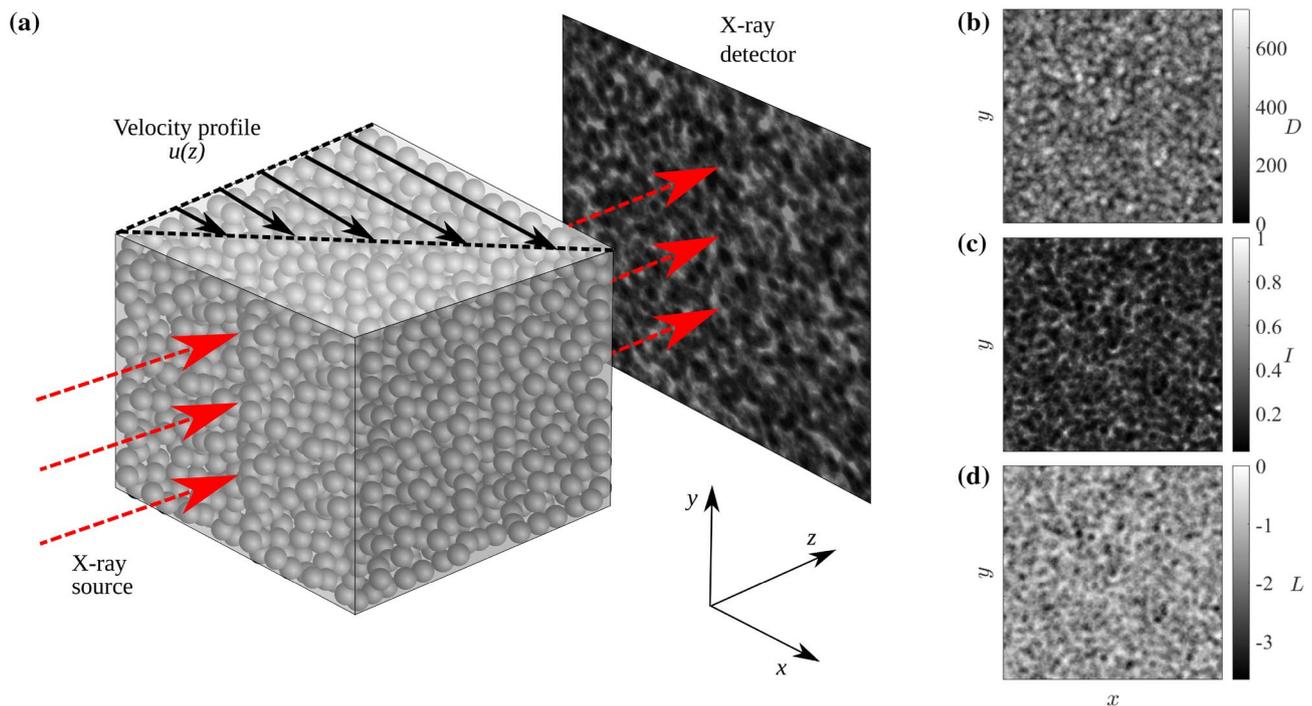


Fig. 9 Schematic plot of the simulation setup for the granular flow x-ray example showing **a** the simulation box filled with spherical particles, the imposed velocity profile $u(z)$ from (35) and projection

direction (red arrows). Panels **b**, **c** and **d** show the projected integrated depth D , intensity I and log of intensity L , respectively

projected depths, intensities and log of intensities (D_2 , I_2 and L_2 , respectively) are computed. This process is repeated over multiple timesteps, resulting in a series of projections of the flowing grains. A movie of these projections is available as supplementary material.

4.3 One-dimensional deconvolution

The next step is, for each pair of projections, to solve a similar deconvolution problem as the previous one-dimensional test case to recover the probability mass function $p(v)$ corresponding to the displacement field $u(z)$. The situation is slightly different here, because the images are now two-dimensional. However, the imposed velocity field, and hence velocity PMF, is still one-dimensional. This means that a one-dimensional deconvolution problem can again be solved and therefore, for simplicity, we proceed by treating each row of pixels individually. For each y position this then gives a pair of one-dimensional signals (varying in the x direction). These are deconvolved in exactly the same way as Sect. 3.3 to give a PMF for each cell, and the final PMF is taken as the average of all such y values. As before, this process is repeated for numerous pairs of images, but is now computed using intensity pairs I_1 , I_2 and their logarithms L_1 , L_2 , as well as the original D_1 , D_2 for completeness.

Figure 10 shows example PMFs from this deconvolution process. It can be seen the average deconvolution is in good agreement with the exact imposed distribution, confirming that deep velocimetry is still capable of recovering the velocity distribution for this more realistic imaging configuration. This is true for all the deconvolution methods (D , I and L). In fact, the reconstructed PMFs from the integrated depths D are identical to those reconstructed from the log of the intensity, L . This is because Eq. (33) indicates that L is simply a linear transformation of D , which does not change the results when computing the deconvolution with a linear

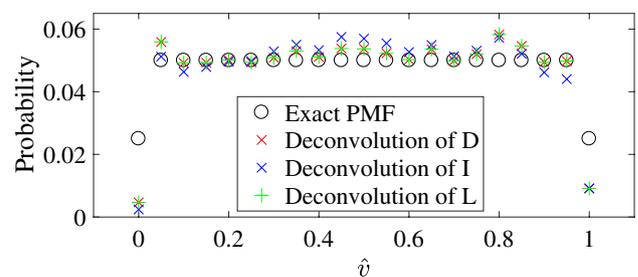


Fig. 10 Plots of the exact underlying PMF for the x-ray system (black circles), example reconstructed PMF from deconvolving the integrated densities D (red crosses), from deconvolving the intensity I (blue crosses) and from deconvolving the log of the intensity L (green pluses). All results represent average of 100 deconvolutions

solver. Using image intensities I , on the other hand, does give different final deconvolutions.

The exact extent of the differences between deconvolutions using D and those using I is not immediately obvious from the single realisations on Fig. 10. For this reason, Fig. 11 explores the deconvolution process in a quantitative manner by computing different numbers of deconvolution averages N_{aves} for each field and plotting the error (26) between the reconstructions and analytical distribution. It is now clear that using the depths D directly (or equivalently the L 's) does produce more accurate results, although the errors are still satisfactorily low when deconvolving the intensities directly and hence the assumptions made in the method derivation are valid. Provisional simulations indicate that the errors are larger when using a higher attenuation coefficient μ_m , since the relationship (31) can no longer be approximated with a linear law. This remains to be explored in full in future research. As for the one-dimensional test case, Fig. 11 shows that the errors consistently decrease with increasing degree of averaging, although the relationship between E and N_{aves} no longer follows a simple scaling law. The errors for this x-ray example are also slightly larger than the first example, which could be due to differences in the sample geometry, solid packing fraction, or imposed velocity profile and should be investigated in more detail in the future. Nevertheless, deep velocimetry produces very promising results for this x-ray imaging system.

4.4 Two-dimensional deconvolution outlook

The row-by-row deconvolution approach employed to recover the one-dimensional velocity distribution from two-dimensional images was chosen here for its simplicity and consistency with the test case of Sect. 3. Alternative methods could be used when handling two-dimensional images that

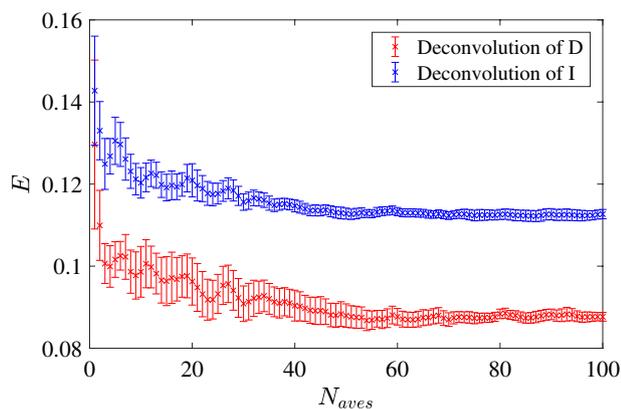


Fig. 11 Plots of the error (26) as a function of number of deconvolution averages for deconvolutions of D or equivalently L (red), and I with $\mu_m = 0.005\text{px}^{-1}$ (blue). Error bars are calculated from computing 7 different results for each point and taking the standard deviation

may, depending on the specific setup, give improved accuracy and/or computational speed. For example, when the motion remains in a single direction, it is possible to reformulate the collection of row-by-row Eqs. (22) into a single matrix equation that takes into account all of the spatial positions simultaneously, which can then be solved in one step. In such an equation, the right-hand-side vector \mathbf{b} would grow from length $(N_x - N_v + 1) + 1$ to $N_y(N_x - N_v + 1) + 1$, where N_y is the number of distinct image rows. Similarly, A would necessarily grow from a $((N_x - N_v + 1) + 1) \times N_v$ matrix to $(N_y(N_x - N_v + 1) + 1) \times N_v$. The case-by-case imaging and computational specifications will likely determine whether this reformulation is favourable over treating each row independently for one-dimensional motion.

In most real flowing systems, the motion will actually be in both planar directions x and y (and possibly in the out-of-plane direction z , but this is not distinguishable from a single projection direction). In this case, the probability density function of the velocity will also have components in the two planar directions, and a more complicated approach will be required to recover this two-dimensional array. The equivalent two-dimensional convolution equation can again be reformulated as a single linear system similar to (22), but the corresponding dimensions become prohibitively large. Similarly, computing the 2D deconvolution by means of constrained nonlinear optimisation is costly due to the typically large number of unknowns in the two-dimensional PDF.

Correlation-based methods are able to overcome these difficulties by treating the two velocity components in isolation. By first computing two-dimensional correlation functions and then averaging these over each spatial dimension in turn, the problem reduces to solving a pair of one-dimensional deconvolutions for each velocity component. The net result is an efficient means of computing the two marginal probability distributions, as opposed to the full joint distribution. A similar dimensionality-reduction approach was tested with the direct deconvolution approach used in this paper. However, it was not found to produce satisfactory results. Analogously to the pre-averaging (in time) described in Sect. 3.4, we attribute this to the fact that spatial averaging homogenises the texture of the images, reducing the accuracy of the deconvolution. Work is ongoing to develop an efficient and accurate method of computing the deconvolution for two-dimensional flows so that deep velocimetry can be applied to more realistic flows.

5 Experimental validation

Having illustrated the potential of the method using artificially generated images, we now proceed to test its capabilities using real experimental data. The experiments involve x-ray radiography of the controlled motion of granular

materials, and therefore complement the simulated results of Sect. 4.

The experimental setup is shown on Fig. 12a. It consists of an acrylic cylinder of inner diameter 79mm, which is filled with glass beads of approximate diameter 3mm. The cylinder is placed on a turntable, allowing it to be rotated about its central (y) axis at a constant rate of 0.56 revolutions per minute. An x-ray source (Spellman XRV generator with Varian NDI-225-21 stationary anode tube) is positioned at a horizontal distance of 1.5 m from the centre of the cylinder, and a detector panel (PaxScan 2520DX) is placed at 0.4 m from the cylinder's centre, in-line with the path of the beam (z -direction). The x-ray source is set to emit radiation at a maximum energy of 180 keV and intensity of 5 mA, and the detector records $1920 \text{ px} \times 1536 \text{ px}$ radiographic images at a rate of 10 frames per second (fps). For our system, the x-ray source emits a continuous stream of radiation which is continuously measured by the detector panel, meaning that the image exposure time is also 10 fps. A movie of these projections is available as supplementary material.

This experimental setup corresponds to solid body rotation of the granular ensemble, with the motion of the grains in the horizontal in-plane (x) and out-of-plane (z) directions. The x-ray beam can be considered to be approximately planar, due to the large ratio between the source-sample and sample-detector distances. The z motion is therefore impossible to determine and we instead focus on the motion in the x direction. This depends on both the x and z positions, relative to the centre of the cylinder. At each x position, the

velocity profile $u_x(\hat{z})$ through the depth of the sample is given by the linear function

$$u_x(\hat{z}) = -u_{max}(x) + 2u_{max}(x)\hat{z}, \quad (36)$$

where the maximum magnitude $u_{max}(x)$ depends on both the rotation rate and position x , and $\hat{z} \in [0, 1]$ represents the rescaled distance from the front to back walls of the cylinder. For the subsequent analysis, we focus on a two-dimensional interrogation window of width 100 px in the x direction and height 400 px in the y direction, centred on the sample midpoint. Variations in the velocity profile (36) with x are very small in this region, and hence the window can be assumed to have a single representative velocity PDF, analogous to those used in Sect. 4 corresponding to the linear profiles (35). Since the in-plane motion is one-dimensional (in the x -direction), the velocity PDF is also one-dimensional and can be reconstructed in a similar method to Sect. 4.3.

Next, we note that the assumptions in deriving the simplified Beer-Lambert law (30) do not strictly apply for this experimental setup. It was assumed that all the material is either grains or interstitial air, and also that the grains have a unique mass attenuation coefficient μ_m . In practise, the path of the x-ray beam is also restricted by the acrylic cylinder, which has an attenuation coefficient that is lower than the grains but is still non-negligible. Furthermore, the source used is polychromatic and emits x-rays at a range of different energies. Because a material's attenuation coefficient depends on this x-ray energy, different energy beams will be

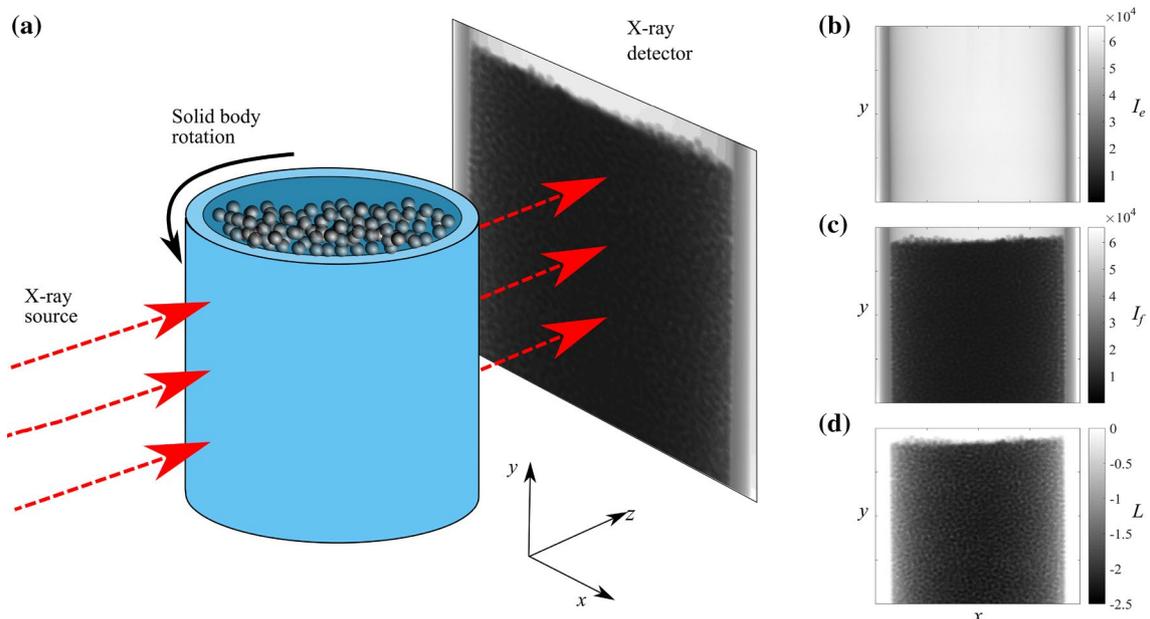


Fig. 12 Schematic plot of the experimental x-ray setup for the controlled rotation of granular material showing a) the cylindrical container filled with 3mm diameter glass beads and projection direction

(red arrows), b) the measured x-ray intensity I_e of the empty container, c) the measured x-ray intensity I_f of the full container and d) the normalised value $L = \log(I_f/I_e)$ used for the deconvolution

absorbed at different rates—a phenomenon known as beam hardening. As a result of these two factors, for our system there is a complex nonlinear relationship between the measured intensity on the detector panel and the depth of moving granular material that the beam has passed through, and we cannot directly apply the methods of Sect. 4. To resolve this complication, x-ray radiographs I_e of the rotating cylinder are also recorded when it is empty of granular material (Fig. 12b). These are then synchronised with the corresponding radiographs of the full container I_f (Fig. 12c). At each timestep, a new intensity image $L := \log(I_f/I_e)$ is computed (Fig. 12d), which explicitly corrects for the rotating cylinder. For the range of interest, L is now found to be well-approximated by a linear function of the depth of granular material D , and thus can be used for subsequent analysis.

The same row-by-row one-dimensional deconvolution process is then applied to these normalised images, and Fig. 13 shows an example probability mass function reconstructed as the average of 50 individual deconvolutions. It can be seen that deep velocimetry is reasonably successful at predicting the full velocity distribution from experimental images, in particular capturing the mean and range of velocities with nonzero probabilities. There are, however, some key differences between the theoretical distributions (black circles on Fig. 13) and the reconstructions (red crosses). Whereas the theoretical distribution gives near-equal weightings to all nonzero-probability displacements, the reconstructed values have a clear maximum at the mean. There is also a degree of smoothing in the reconstructed values that is not present in the theoretical profile.

Figure 14 quantifies these discrepancies further, plotting the deconvolution error (26) as a function of the number of averages, N_{aves} . As found previously, there is an initial decrease in the error with more runs, but this saturates at around 16% – approximately twice the value measured for the theoretical results in Sect. 4. It is not particularly surprising that the experimental results give higher errors than the simulated radiographs. Indeed, experiments are inherently noisy, with the quality of the resulting images being

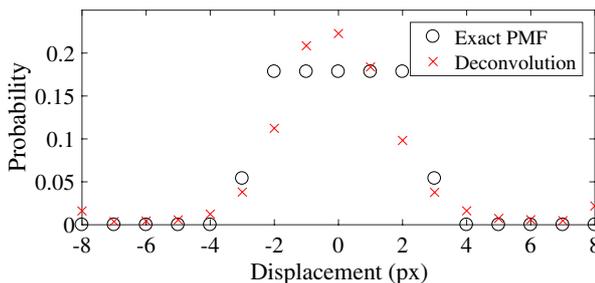


Fig. 13 Plots of the exact underlying PMF for the x-ray rotation experiments (black circles), and the corresponding deconvolutions (red crosses) computed by averaging 50 computations

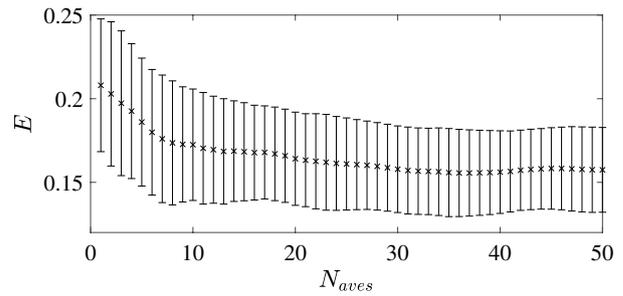


Fig. 14 Plots of the error (26) as a function of number of deconvolution averages. Error bars are calculated from computing 20 different results for each point and taking the standard deviation

degraded by the non-uniform intensity of the x-ray source, the scattering and diffraction of x-rays, and the finite radiograph resolution. These factors were not taken into consideration for the simulated x-ray images. Furthermore, the process of beam-hardening was not modelled in the simulated images, whereas this does occur in the experiments. The normalisation process applied to the experimental radiographs was designed to recover an approximately linear relationship between intensity and depth, but remaining variations could also be a contributing factor to the reconstruction errors. Nevertheless, deep velocimetry on these experimental images is still able to capture the key motions, and we thus believe that it offers great promise as an experimental technique. In the future it may also be possible to precisely quantify the effect of experimental noise and beam-hardening on the reconstruction process, and thus explicitly correct the resulting distributions, but this correction is beyond the scope of this introductory paper.

6 Other applications

The choice of imaging system in Sects. 4 and 5 could be of immediate interest to granular materials researchers with existing x-ray radiography capabilities. While the use of x-ray CT to study the internal structure of such opaque systems is now widespread, this is typically restricted to quasi-static deformations. X-ray CT therefore offers limited insight into the dynamics of continuously flowing granular media. However, the same hardware could easily be repurposed for deep velocimetry by recording successive x-ray radiographs from a fixed direction, as opposed to continuously rotating the sample. In this way, deep velocimetry using x-rays offers researchers a means of measuring velocity distributions in a range of dense granular flows, which have applications spanning the natural (e.g. snow avalanches, debris flows, landslides) and man-made (e.g. mineral processing, pharmaceuticals, food processing) environments.

Furthermore, dynamic x-ray radiography is a general projected imaging method that is not limited to dense granular flows. We therefore anticipate that a similar approach could also be adopted by researchers working with other deformable media, as long as the media has a sufficiently heterogeneous density profile to provide the required radiograph texture. Such materials might include foams, gels, suspensions and different forms of soils and rocks. Of particular interest is the application of deep velocimetry to successive x-ray radiographs of viscous fluids such as water - fluids that are typically investigated using regular PIV with optical imaging. In this case, tracer particles with different x-ray attenuation coefficients to the bulk fluid can be inserted into the flow, which then provide the appropriate radiograph contrast for image analysis. It has previously been demonstrated that regular PIV can be applied to such radiographs (e.g. Lee and Kim 2003) to measure the single most likely velocity (through the path of the x-ray beam) in each patch. Alternatively, some complex non-Newtonian fluids like blood already have discernible density fluctuations that make their radiographs amenable to PIV (e.g. Kim and Lee 2006), and in other cases CO₂ microbubbles can also provide appropriate tracer texture (Park et al. 2015). In all of these scenarios, deep velocimetry, as opposed to PIV, could be applied to the same radiographs, which would provide valuable additional out-of-plane velocity distribution information.

Going beyond x-ray radiography, deep velocimetry could also be applied to other projection-based imaging systems. While the ideal system for the theory of Sect. 2 is one that measures the integrated density field D directly, the method could also work for other indirect, proxy measures of integrated density. This is especially true if the measured intensity I is a linear function of D , or if it can be closely approximated by a linear function. Even in highly nonlinear cases, deep velocimetry could be applied as long as there is a well-understood, one-to-one relationship between I and D , since this allows the measured I to be converted back to work with D directly.

Other projection-based imaging systems potentially amenable to deep velocimetry include shadowgraphy, and the closely related Schlieren imaging (Luthman et al. 2019; Ozawa et al. 2020). These both exploit the fact that fluid regions of different densities have different refractive indexes, and produce projected images representative of a finite depth of flowing material. The projected images typically have much sharper contrast than is visible with the human eye, and can therefore be used to measure flow patterns over this integrated depth. Deep velocimetry could again be applied directly to the same images to provide additional and complementary measurements. Finally, it may also be possible to work with regular optical images of seeded flows of transparent fluids that look through a finite flow depth, as opposed to on a single plane of a laser. In

this case a final image is built up by summing all opaque tracer particles that obstruct and cast a shadow on the screen. Such projected images could certainly be amenable to this new approach. Deep velocimetry could therefore provide new insight into a range of finite-depth systems, for example droplet dynamics (Gultekin et al. 2020), that have only previously been examined on a single imaging plane.

7 Conclusions

This paper has presented a new image analysis technique for measuring velocities using projection-based imaging systems. Building on the established particle image velocimetry (PIV) method, where image correlation is used to recover a single velocity value in each interrogation window, we exploit the extra out-of-plane information provided by projection imaging to reconstruct the full velocity distribution through the imaging direction. Crucially, the new method, which we call deep velocimetry, does not require computation of any correlation functions. Instead it involves deconvolving a pair of sequential images directly to recover the velocity distributions.

In Sect. 2 the theory behind deep velocimetry has been described mathematically, indicating the important assumptions that are required for the method to work on images of integrated density fields. This theory is tested for a simple one-dimensional test case in Sect. 3, where we demonstrate that deconvolving such integrated density fields is a highly accurate method of recovering the full distribution of velocities through the projection direction. Key control parameters are found to be the solids volume fraction, as well as depth of integrated material, with errors decreasing as both are increased. We also compare the method to previous correlation-based approaches, showing that, for the one-dimensional test case, it is both more accurate and less computationally intensive. The new method is also fairly robust to the addition of image noise, again performing equally as well as previous methods up until the point where all approaches break down.

Section 4 then demonstrates how deep velocimetry can be applied to a more realistic projection-based imaging system, specifically x-ray radiography. In this case the x-ray detector measures the incident x-ray intensity, not the integrated material density directly. We show that deconvolving these intensity fields, which are a nonlinear function of the integrated density D , again provides accurate reconstructions of the full velocity distributions, although some additional errors are introduced due to this nonlinearity. These additional errors can be avoided by taking logarithms L of the intensity fields before carrying out the appropriate deconvolution. Since L is a linear function of D , the deconvolution process yields identical

results as deconvolving D directly. In all cases, a full study of the important parameters controlling the errors would be a worthwhile future exercise to establish the limitations and best practise for deep velocimetry applied to x-ray imaging.

The simulated test cases investigated in this paper all restrict the in-plane motion to one direction, whereas most real flows of interest will move in two in-plane directions. An important extension of the method is therefore to generalise the reconstructive procedure to account for general planar motion. This remains an ongoing exercise due to the additional computational complexity that previous correlation-based methods do not face in 2D, as described in Sect. 4.4. Nevertheless, the promising results for deep velocimetry of one-dimensional motion strongly suggest that the extension is achievable and worth pursuing.

The method is tested experimentally in Sect. 5 using the controlled rotation of granular materials. Deep velocimetry is able to reconstruct the key features of this motion, with the measured errors approximately twice those found for the equivalent artificial results. The corresponding distributions display some additional smoothing, which could be due to the inherent noise of the experimental measurements as well as beam-hardening effects. It may be possible to quantify and explicitly correct for such factors to further improve the method in the future.

Deep velocimetry could also be applied to other projection-based imaging systems, which is discussed in Sect. 6. In all applications, the added value of full velocity distributions, as opposed to the single modal quantities from classical PIV, is important, especially in systems with strong gradients in the out-of-plane direction. Furthermore, whilst a single projection direction does not give any information about where each velocity occurs in the out-of-plane direction, it may be possible to combine deep velocimetry with additional flow assumptions, for example by considering axisymmetric systems (e.g. Pimenta et al. 2020), or with simultaneous imaging directions (as in x-ray rheography (Baker et al. 2018)) to reconstruct fully three-dimensional fields.

Acknowledgements The authors would like to thank Benjy Marks and François Guillard for interesting discussions.

Author contributions James Lindsay Baker developed the theory, conducted the simulations and experiments and wrote the paper. Itai Einav oversaw the project and ideas, and edited the manuscript.

Funding This research was supported by a Discovery Early Career Researcher Award from the Australian Research Council (DE200101116).

Data availability The code and data used in this paper are available from the corresponding author on request.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Baker J, Guillard F, Marks B, Einav I (2018) X-ray rheography uncovers planar granular flows despite non-planar walls. *Nat Commun* 9(1):5119 (ISSN 2041-1723)
- Bhattacharya S, Vlachos PP (2020) Volumetric particle tracking velocimetry (PTV) uncertainty quantification. *Exp Fluids* 61(9):1–18. <https://doi.org/10.1007/s00348-020-03021-6> (ISSN 14321114)
- Brücker C (1997) 3D scanning PIV applied to an air flow in a motored engine using digital high-speed video. *Meas Sci Technol* 8(12):1480–1492. <https://doi.org/10.1088/0957-0233/8/12/011> (ISSN 09570233)
- Delnoij E, Westerweel J, Deen NG, Kuipers JA, Van Swaaij WP (1999) Ensemble correlation PIV applied to bubble plumes rising in a bubble column. *Chem Eng Sci* 54(21):5159–5171. [https://doi.org/10.1016/S0009-2509\(99\)00233-X](https://doi.org/10.1016/S0009-2509(99)00233-X) (ISSN 00092509)
- Dijksman JA, Rietz F, Lorincz KA, Van Hecke M, Losert W (2012) Invited Article: Refractive index matched scanning of dense granular materials. *Rev. Sci. Instrum.* 83(1). <https://doi.org/10.1063/1.3674173> (ISSN 00346748)
- Dubsky S, Jamison RA, Irvine SC, Siu KKW, Hourigan K, Fouras A (2010) Computed tomographic x-ray velocimetry. *Appl Phys Lett* 96(2):023702. <https://doi.org/10.1063/1.3285173> (ISSN 00036951)
- Elsinga GE, Scarano F, Wieneke B, Van Oudheusden BW (2006) Tomographic particle image velocimetry. *Exp Fluids* 41(6):933–947. <https://doi.org/10.1007/s00348-006-0212-z> (ISSN 07234864)
- Fouras A, Dusting J, Lewis R, Hourigan K (2007) Three-dimensional synchrotron x-ray particle image velocimetry. *J. Appl. Phys.* 102(6). <https://doi.org/10.1063/1.2783978> (ISSN 00218979)
- Guillard F, Marks B, Einav I (2017) Dynamic X-ray radiography reveals particle size and shape orientation fields during granular flow. *Sci Rep* 40:1–13 (ISSN 2045-2322)
- Gultekin A, Erkan N, Colak U, Suzuki S (2020) PIV measurement inside single and double droplet interaction on a solid surface. *Exp. Fluids* 1–18. <https://doi.org/10.1007/s00348-020-03051-0> (ISSN 1432-1114)
- Hall S, Bornert M, Desrues J, Pannier Y, Lenoir N, Viggiani G, Bésuelle P (2010) Discrete and continuum analysis of localised deformation in sand using X-ray μ CT and volumetric digital image correlation. *Géotechnique* 60(5):315–322 (ISSN 0016-8505)

- Kähler CJ, Kompenhans J (2000) Fundamentals of multiple plane stereo particle image velocimetry. *Exp Fluids* 29(SUPPL. 1):70–77. <https://doi.org/10.1007/s003480070009> (ISSN 07234864)
- Katz J, Sheng J (2010) Applications of Holography in Fluid Mechanics and Particle Dynamics. *Annu Rev Fluid Mech* 42(1):531–555. <https://doi.org/10.1146/annurev-fluid-121108-145508> (ISSN 0066-4189)
- Keane RD, Adrian RJ (1992) Theory of cross-correlation analysis of PIV images. *Appl Sci Res* 49(3):191–215. <https://doi.org/10.1007/BF00384623> (ISSN 00036994)
- Kim GB, Lee SJ (2006) X-ray PIV measurements of blood flows without tracer particles. *Exp Fluids* 41(2):195–200. <https://doi.org/10.1007/s00348-006-0147-4> (ISSN 07234864)
- Lee SJ, Kim GB (2003) X-ray particle image velocimetry for measuring quantitative flow information inside opaque objects. *J Appl Phys* 94(5):3620–3623 (ISSN 00218979)
- Lueptow RM, Akonur A, Shinbrot T (2000) PIV for granular flows. *Exp Fluids* 28(2):183–186. <https://doi.org/10.1007/s00348-020-03021-60> (ISSN 0723-4864)
- Luthman E, Cymbalist N, Lang D, Candler G, Dimotakis P (2019) Simulating schlieren and shadowgraph images from LES data. *Exp Fluids* 60(8):1–16. <https://doi.org/10.1007/s00348-020-03021-61> (ISSN 14321114)
- Meinhart CD, Wereley ST, Santiago JG (2000) A piv algorithm for estimating time-averaged velocity fields. *J Fluids Eng Trans ASME* 122(2):285–289. <https://doi.org/10.1007/s00348-020-03021-62> (ISSN 1528901X)
- Ozawa Y, Ibuki T, Nonomura T, Suzuki K, Komuro A, Ando A, Asai K (2020) Single-pixel resolution velocity/convection velocity field of a supersonic jet measured by particle/schlieren image velocimetry. *Exp Fluids* 61(6):1–18. <https://doi.org/10.1007/s00348-020-03021-63> (ISSN 14321114)
- Park H, Yeom E, Seo S-J, Lim J-H, Lee S-J (2015) Measurement of real pulsatile blood flow using X-ray PIV technique with CO₂ microbubbles. *Sci Rep* 5:8840. <https://doi.org/10.1007/s00348-020-03021-64> (ISSN 2045-2322)
- Parker DJ (2017) Positron emission particle tracking and its application to granular media. *Rev Sci Instrum* 88:051803. <https://doi.org/10.1007/s00348-020-03021-65>
- Pimenta F, Toda-Peters K, Shen AQ, Alves MA, Haward SJ (2020) Viscous flow through microfabricated axisymmetric contraction/expansion geometries. *Exp Fluids* 61(9):1–16. <https://doi.org/10.1007/s00348-020-03021-66> (ISSN 14321114)
- Prasad AK (2000) Stereoscopic particle image velocimetry. *Exp Fluids* 29(2):103–116. <https://doi.org/10.1007/s00348-020-03021-67> (ISSN 07234864)
- Qiao Z, Wang Z, Zhang C, Yuan S, Zhu Y, Wang J (2007) Extension of PIV for Measuring Granular Temperature Field in Dense Fluidized Beds. *AIChE J* 53(1):108–118. <https://doi.org/10.1007/s00348-020-03021-68> (ISSN 12350621)
- Sarno L, Carravetta A, Tai YC, Martino R, Papa MN, Kuo CY. Measuring the velocity fields of granular flows - Employment of a multi-pass two-dimensional particle image velocimetry (2D-PIV) approach. *Adv. Powder Technol.*, (August), 2018. ISSN 09218831. <https://doi.org/10.1016/j.appt.2018.08.014>
- Scharnowski S, Hain R, Kähler CJ (2012) Reynolds stress estimation up to single-pixel resolution using PIV-measurements. *Exp Fluids* 52(4):985–1002. <https://doi.org/10.1007/s00348-020-03021-69> (ISSN 07234864)
- Stannarius R (2017) Magnetic Resonance Imaging of Granular Materials. *Rev Sci Instrum* 88:051806. <https://doi.org/10.1088/0957-0233/8/12/0110> (ISSN 0034-6748)
- Westerweel J (2008) On velocity gradients in PIV interrogation. *Exp Fluids* 44(5):831–842. <https://doi.org/10.1088/0957-0233/8/12/0111> (ISSN 07234864)
- Westerweel J, Elsinga GE, Adrian RJ (2013) Particle Image Velocimetry for Complex and Turbulent Flows. *Annu Rev Fluid Mech* 45(1):409–436. <https://doi.org/10.1088/0957-0233/8/12/0112> (ISSN 0066-4189)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.