

Modified control grid interpolation for the volumetric reconstruction of fluid flows

David H. Frakes · Kerem Pekkan · Lakshmi P. Dasi · Hiroumi D. Kitajima ·
Diane de Zelicourt · Hwa Liang Leo · Josie Carberry · Kartik Sundareswaran ·
Helene Simon · Ajit P. Yoganathan

Received: 4 August 2007 / Revised: 12 April 2008 / Accepted: 11 May 2008 / Published online: 13 June 2008
© Springer-Verlag 2008

Abstract Complex applications in fluid dynamics research often require more highly resolved velocity data than direct measurements or simulations provide. The advent of stereo PIV and PCMR techniques has advanced the state-of-the-art in flow velocity measurement, but 3D spatial resolution remains limited. Here a new technique is proposed for velocity data interpolation to address this problem. The new method performs with higher quality than competing solutions from the literature in terms of accurately interpolating velocities, maintaining fluid structure and domain boundaries, and preserving coherent structures.

1 Introduction

Fluids experimentation routinely involves the sampling of flow domains with measurement planes. Applications that

analyze three-dimensional (3D) data often demand velocity quantification at points other than discrete measurement locations. In these cases interpolation is required. When points of interest lie between measurement planes, interpolation calculations may rely on relatively distant samples. Even when planes are acquired as densely as experimentally possible, they may still be separated by far greater distances than in-plane measurements. This is true for both phase contrast magnetic resonance (PCMR) data, and for particle image velocimetry (PIV) data (Sakakibara et al. 2004; Ozturk et al. 2003).

Inter-plane interpolation problems in the contexts of PIV and magnetic resonance (MR) imaging are very similar. Both modalities acquire velocity data in planes (images) and are quantitative. Previous image processing solutions for general inter-plane interpolation problems fall into two categories, scene-based and object-based (Penney et al. 2004). The scene-based variety, one example of which is cubic spline interpolation, considers only measurement point values and their coordinates in interpolating, while object-based methods exploit more complex image information (Grevera and Udupa 1996). Scene-based solutions are commonly used for in-plane problems, where measurements are close to the interpolation site. They perform less impressively for interpolation between planes. Object-based techniques have long been considered superior in terms of quality for the latter case. Nevertheless, scene-based methods remain popular for inter-plane velocity interpolation (Kimura and Herring 1996; Yeung and Pope 1988; Karatekin et al. 1998). This may be because most object-based methods were designed for, and have been applied primarily to, anatomical data. In this paper a novel object-based strategy is presented for interpolating incompressible velocity data. The new methodology is founded in medical imaging, but has been optimized for

D. H. Frakes (✉)
Harrington Department of Bioengineering,
Arizona State University, Tempe, AZ, USA
e-mail: dfrakes@asu.edu

D. H. Frakes
4-D Imaging, Inc., Atlanta, GA, USA

K. Pekkan
Department of Biomedical Engineering,
Carnegie Mellon University, Pittsburgh, PA, USA

L. P. Dasi · H. D. Kitajima · D. de Zelicourt ·
H. L. Leo · J. Carberry · K. Sundareswaran · H. Simon ·
A. P. Yoganathan
Wallace H. Coulter Department of Biomedical Engineering,
Georgia Institute of Technology, Atlanta, GA, USA

performance in the context of fluid flows. It is demonstrated for experimental bio-fluids applications, but translates well to other multi-dimensional divergence-free vector field reconstruction problems.

This paper describes data that span 3D volumes and that comprise vectors with multiple components. For clarity, the term 3D is used here to refer to the spatial nature of data rather than the number of components in a vector. When necessary, vectors are described explicitly as having a specific number of components.

1.1 Motivation

Experimental fluid dynamic applications span a broad range of active research fields including many branches of physics and engineering. Many of these use interpolation to approximate velocity values at points of interest where direct measurements are unavailable (Karatekin et al. 1998; Derou et al. 1995; Elkins and Alley 2007). Two examples of such applications are oblique plane visualization and numerical particle tracking. Both underscore the importance of interpolation for success. Consider the example where a flow field needs to be visualized in an orientation that is orthogonal to a stack of acquired measurement planes. Most of the values composing a plane of uniformly resolved data in this orientation will be interpolated. Accordingly, accurate representation of such a visualization plane depends heavily on interpolator performance. For numerical particle tracking (the determination of particle trajectories integrated through experimentally acquired or simulated velocity fields) the implications are similar. Calculated particle trajectories will inevitably pass through points where velocities have not been measured directly or where simulated values are undefined, even for very finely resolved data. Clearly the interpolated velocities used for path integration will have a dramatic impact on accuracy. When measurement planes are acquired sparsely, the implications of interpolator choice are especially significant. This is also the case for exercises involving spatial derivative calculations. However, even when acquired or simulated data are isotropically resolved, interpolating between measurement points is analogous to an inter-plane interpolation problem where in-plane resolution and plane spacing are equal. The methodology proposed here provides an improved solution to the velocity interpolation problem for a wide range of applications that rely on 3D data from incompressible flows.

1.2 Interpolation

Conventional approaches to velocity data interpolation include many scene-based interpolators that approximate a sinc function. For example, cubic spline interpolation has

been popular in recent literature for numerical particle tracking (Melnikov and Shevtsova 2005; Yeung 2002). The problem with this type of interpolator is not a function of the interpolator itself, but rather the data points it operates on. Figure 1 indicates two measurement planes that intersect a pipe flow. The red squares occupy the same location within their respective planes. Linearly interpolated values between these two points would be based on one point within the fluid domain and one outside of it. This circumstance is not unrealistic as standard PIV measurements from wall-bounded flows often include the boundary. A more complex scene-based interpolator would incorporate additional samples from more distant planes, but the data points being operated upon would still be fundamentally unrelated. The case at the boundary is extreme and can be handled uniquely to improve results. For example, masking techniques can be applied to isolate the fluid domain prior to interpolating. These can perform unreliably, however, and are impractical for large data sets. Regardless, the same general problem persists throughout the fluid domain: points at the same location in different planes often belong to relatively unrelated fluid structures.

This shortcoming of scene-based interpolators has been recognized in the medical imaging arena, and numerous object-based directional interpolators have been proposed to address it for anatomical data. These interpolators operate along intelligently determined vectors that can alleviate the general mismatching problem outlined in Fig. 1. Shape, morphology, and registration-based interpolators are several examples (Penney et al. 2004; Grevera




Fig. 1 This example illustrates the implications of using scene-based methods in the context of velocity data interpolation. Two acquisition planes sampling fluid flow through a pipe are indicated by the translucent *gray traces*. The *red squares* indicate points that occupy the same location within respective planes. A purely scene-based interpolator would make use of these two points, one inside of the fluid domain and one outside of it, to interpolate velocity values at points in between

and Udupa 1996). Adaptive control grid interpolation (ACGI), a form of which is proposed here for velocity interpolation, is another (Frakes et al. 2003). In the context of interpolating velocity data, most of these techniques suffer due to a limited region of support. They consider only the two immediate measurement planes bounding the interpolation site. Furthermore, some of these techniques accomplish interpolation only between certain regions of neighboring images. For applications that demand the reconstruction of a global fluid domain or interpolation at arbitrary coordinates, such as numerical particle tracking, such results would be inadequate. Lastly these techniques were designed for and have been applied primarily to anatomical data. The velocity interpolation problem is different in that well-defined physical laws govern the spatial variation of velocity within a fluid domain as opposed to the spatial variation of anatomy within an imaging subject.

ACGI is built upon optical flow and produces displacement fields linking each image completely to its neighbors. These links, or registrations, allow similarities between related fluid structures at different locations in different images to be recognized. Previous research on MR velocity data interpolation has demonstrated the ability of simple linear ACGI to outperform traditional means in a fluids context (Frakes et al. 2004). The proposed technique differs from previous methods in several ways. It combines adjacent displacement fields to enable directional interpolation with an extended region of support. Fluids-specific divergence information is also integrated to facilitate high-quality performance in the context of incompressible flow data.

2 Methods

The new algorithm is built upon ACGI, which is outlined next as background. ACGI is rooted in image processing, so the foundations of the new method are described as they apply to interpolating a scalar field. The transition to vector fields is covered later as one modification that defines the new approach.

2.1 ACGI

Motion estimation techniques allow similar features from different images to be registered. They provide a displacement field that maps points in one image to points in another so as to connect those features. ACGI is a hybrid motion estimator that incorporates components of both block-matching and optical flow. It has been used extensively for video coding, target tracking, and data reconstruction (Frakes et al. 2003; Monaco 1997; Frakes

et al. 2001). Like block-matching, which directly compares regions from one image to regions from another in order to identify highly correlated structures, ACGI partitions the image into rectangular sub-regions. However, ACGI differs in that it determines a unique displacement for each point within an image based on the optical flow constraint equation (OFCE), which can be written as

$$I[n_1, n_2, k] = I(n_1 + d_1[n_1, n_2, k], n_2 + d_2[n_1, n_2, k], k + \delta k), \quad (1)$$

where $I[n_1, n_2, k]$ represents the scalar value for a given point at coordinates (n_1, n_2) within image k , and $d_1[n_1, n_2, k]$ and $d_2[n_1, n_2, k]$ represent the displacement of that point (Horn 1968). A unique solution to Eq. 1 is found by assuming displacement field smoothness, a constraint that is imposed here through a connected bilinear motion model.

Optical flow-based methods rely on the assumption that scalar values remain constant as they transition from one plane to the next. Since this assumption is usually flawed, the error associated with the OFCE will be non-zero. However, minimizing that error still enables high-quality interpolation, as the results to follow indicate (Frakes et al. 2008).

In the ACGI formulation, displacement vectors for points within a sub-region are related to the displacements for the four points bounding that region, control points, through bilinear interpolation. Within this framework the OFCE error is minimized. Expressing displacements in vector form, that error becomes:

$$E(\bar{\alpha}, \bar{\beta}) = \sum \sum_{\mathbf{n} \in R} (I[\mathbf{n}, k] - I[n_1 + \bar{\alpha}^T \bar{\theta}(\mathbf{n}), n_2 + \bar{\beta}^T \bar{\phi}(\mathbf{n}), k + \delta k])^2. \quad (2)$$

In this expression, $\bar{\alpha}$ and $\bar{\beta}$ represent vectors composed of the row and column components of control point displacements, the vector $\mathbf{n} = (n_1, n_2)$ denotes coordinates within an image, and $\bar{\theta}$ and $\bar{\phi}$ represent basis functions that implement the bilinear interpolation mentioned earlier. The error is summed over the rectangular sub-region R . Displacement field smoothness is inherently imposed by the bilinear motion model and by model connection, which dictates that the displacements for a control point in one block are the same as the displacements for that point with respect to the other blocks it belongs to. Error minimization is performed iteratively for each block in an image, and the result is a displacement field that is piecewise smooth and globally continuous.

This general formulation describes how a displacement field is calculated. Partitioning images into sub-regions allows for more efficient computation than with pure optical flow. However, the optical flow foundation is still able to characterize the complex displacements, unlike

block-matching, that must be captured for high-quality directional interpolation. This combination achieves efficiency without sacrificing quality. The overviews of optical flow and ACGI presented here are intended to summarize the foundations of this work. More information on these general techniques, including complete descriptions in equation form and implementation details, can be found in the literature (Horn 1968; Frakes et al. 2008; Chen and Willson 2000).

2.2 Modifications

Several significant modifications were made to the traditional ACGI model in order to implement ACGI2. Adjacent displacement field linkage and improved displacement field optimization mechanics both contribute to higher-quality performance in the context of velocity data. Determination of the appropriate optimization framework was carried out with displacement field linkage, described next, already in place.

2.2.1 Displacement field linkage

ACGI allows each image to be linked completely to each of its neighbors. By following a displacement vector from one image to its neighbor, then from that neighbor to its neighbor, multiple planes can be linked. Using this principle, interpolation can be carried out based on data outside of the immediate image pair that bounds an interpolation site. In this work interpolation between images is performed with a fourth-order accurate twice-continuously differentiable cubic spline interpolator, as described by Yeung (2002), using data points from the closest four images to the interpolation site. The appropriate data points to be used in the interpolation calculation are determined based on the aforementioned displacement fields. The result is that a filter more closely approximating the theoretically optimal infinite sinc function can be employed, but the points it is applied to are chosen more judiciously than in the scene-based case.

2.2.2 Optimization framework

With respect to velocity data within a fluid domain, more information is available to guide interpolation than in the context of anatomical medical data. One example is that the divergence for any volume within an incompressible fluid flow must equal zero. This fact was used advantageously in designing ACGI2 as both divergence errors and interpolated velocity value errors were taken into account in identifying the optimal mechanics for the algorithm. As a result the proposed technique is applicable only to incompressible flows. The specifics of how these two

factors were used to implement the algorithm are described in the following paragraphs.

In previous implementations of both ACGI and other optical flow-based interpolation techniques, it has been assumed that the displacement field which links one image to another with minimal OFCE error will lead to the most accurate interpolation. However, empirical research suggests that this is not necessarily the case. Standard ACGI implementations start with a low-resolution grid, corresponding to a large region R as described in Eq. 2, and then proceed to higher resolution grids using the displacement fields from previous iterations as initial conditions, as long as OFCE error is decreased. This is similar to multi-grid optical flow approaches, and typically leads to displacement fields that facilitate higher quality interpolation as both large and small-scale displacements are captured (Zhao and Sawhney 2002). However, in the context of velocity data this framework often results in greater divergence and interpolated velocity value errors than alternatives.

In order to define the ACGI2 optimization mechanics, a database of approximately 4,000 MR and PIV images was considered. For each set of images in this database, every other image was removed and the complete set was reconstructed. Velocity magnitudes from interpolated images, for different optimization frameworks and parameters, were then related to corresponding values from known images via mean squared error (MSE). Future references to interpolation error refer to this quantity. Divergence was evaluated similarly in relation to the theoretical standard (zero). Interpolation quality was judged based on these metrics. Through this process one framework/parameter combination emerged as optimal.

Two salient features of the optimization framework distinguish it from standard ACGI. The first relates accepting initial displacement field updates. The optimization begins using a grid comprised of 64 identical rectangular block sub-regions (coarse grid). These correspond to R in Eq. 2. Fields of view were fairly consistent across the image database used here, making a static upper bound on block size sufficient. However, for images acquired on drastically different scales, adjustments to the coarse grid resolution may be appropriate. If the displacement field calculated for this grid decreases OFCE error by 50% or more, the results are accepted and successively finer grids, by a factor of 2, are applied until the grid resolution cannot be refined any further without treating every pixel as a control point (fine grid). If the 50% decrease is not observed for the coarse grid, the displacement field update is rejected, and the fine grid is applied next with a zero displacement initialization. Using a more highly resolved fine grid led to increased interpolation error in 78.3% of trials and was accompanied by dramatic

increases in computation time. All grid refinements prior to reaching the fine grid were highly beneficial on average. The second differentiating feature relates to the termination criteria for the minimization of OFCE error. If an update for an individual block decreases that block's OFCE error by an amount less than 10%, the update is rejected and the error minimization terminates.

In summary, a 50% OFCE error decrease is required for initial displacement field updates to be accepted, and a 10% decrease is required for subsequent updates to be accepted. These parameters were determined experimentally via examination of the aforementioned image database. The first parameter was varied from 0 to 70% and the second from 0 to 30%. Performance was quantified in terms of interpolation error and divergence error, both of which were minimized at the respective parameter values of 50 and 10%. Clear error trends associated with variations in both parameters were observed. A parallel study conducted on anatomical data also identified these approximate points as optimal, and by implementing the modified framework, achieved an average decrease in interpolation error of 8.80% as compared to standard ACGI.

The ACGI2 formulation is unique since the criteria for termination in previous implementations have focused on either efficiency or OFCE error. The new criteria focus on interpolation and divergence errors. The proposed method integrates fluids-specific information into the velocity interpolation process, enabling an improved solution for the unique problem addressed here.

2.2.3 Velocity component considerations

While it is an obvious modification, ACGI2 in the context of flow data is also differentiated from standard ACGI in that it accounts for interpolating vector as opposed to scalar data. For any pair of measurement planes sampling all three components of velocity, there are three separate images from each plane that must be accounted for in the displacement field calculation: the u , v , and w components of velocity. While there are arguments for deriving unique displacements for each component, poor results were observed for all such implementations that were explored. In particular divergence errors were excessive when each component was interpolated uniquely. Within the framework presented here, the use of velocity magnitude images as the basis for displacement calculations led to better results in terms of both interpolation and divergence errors as compared to using any one of the velocity components alone, or any alternative combination thereof. Fortuitously, this approach also demands less computation than working with each component individually.

3 Results

Three validations are presented to demonstrate the capabilities of the proposed technique. They include both general fundamental cases and several real-world biomedical applications that highlight the utility of the new method. The validations begin with a simple single-component velocity case and advance to two and three-component velocity fields that describe progressively more complex flows. In two of the validations a ground truth solution for the velocity field is available and interpolation error, as previously defined, is used for comparison. Ground truth is established through CFD simulation or other numerical means. Two of the validations also provide all three components of velocity, allowing divergence MSE values to be used for comparison. Lastly, vorticity is examined in the final case where coherent structures are prevalent to demonstrate that ACGI2 can accommodate spatially rich flow. In every validation the new method is compared to a competitor, a fourth-order accurate twice-continuously differentiable cubic spline interpolator as described by Yeung and Pope (1988), Yeung (2002) and adopted by Kimura and Herring (1996). This is the same interpolator used in the ACGI2 framework where it operates on points chosen based on the displacement fields. Linear interpolation is also evaluated in the first test case as in Melnikov and Shevtsova (2005).

3.1 Validation 1

The first validation involves four simulated pipe flows, corresponding to parabolic, skewed parabolic, trigonometric, and exponential flow profiles. For each pipe flow a unique, isotropically resolved cubic volume of 256 samples per side was formed. A cylindrical fluid domain was established within each volume and 8-bit velocity values were assigned at every discrete location within the domain based on the respective governing equations for each type of flow (with a maximum value of 255). Planes were extracted from each volume at five spacings consistent with in-plane to out-of-plane resolution ratios common to both MRI and PIV. An example set of planes is illustrated in Fig. 2 where the measurement plane spacing is near the middle of the range explored. Each set of planes was reconstructed with three different interpolators (linear, cubic spline, and ACGI2) to match the resolution of the originally sampled volume. Reconstructed versions were compared to the originals in order to evaluate interpolation error, which was summed only over the portion of the volume spanned by the samples.

Figure 3 displays a log–log plot of RMS errors for each measurement plane spacing, averaged over the four flow profiles examined. A line was fit to the data for each




Fig. 2 Illustration of a cubic volume bounding a pipe flow where the velocities are assigned to represent a parabolic flow profile. The *red lines* indicate simulated acquisition planes taken at a consistent spacing. The *green line* indicates the center plane of the volume for future reference





Fig. 3 Plots of the RMS error values for linear, cubic spline, and ACGI2 interpolation. These RMS values were averaged for the four flow profiles examined. Linear fits to each data series show that the steepest trendline, and accordingly the highest experimental order of accuracy, corresponds to ACGI2

Fig. 4 Flow profiles reconstructed with various interpolators from samples of a parabolic flow profile. Here the *color map* indicates velocity magnitude and the *white lines* represent isocontours. The reconstruction created with ACGI2 (c) resembles the original (a) most closely and displays fewer artifacts than the cubic spline interpolation (b)



interpolation methodology to illustrate the trend in error increase with respect to plane spacing. The steepest line corresponds to the data points from the ACGI2 series, indicating the highest experimental order of accuracy. The observed experimental orders of accuracy were 3.65, 3.21, and 2.29 for ACGI2, cubic spline, and linear interpolation, respectively. Several trends were consistent across all of the flow profiles. For every flow profile and measurement plane spacing, the errors for ACGI2 were lower than those for the other interpolators. Also, the steepest trendline, indicating the highest experimental order of accuracy, was contributed by the ACGI2 data series in each case.

Strictly as a point of reference, average experimental orders of accuracy were also evaluated for interpolating the center plane of the fluid domain alone. This plane is represented by the green line orthogonal to the measurement planes in Fig. 2. Values of 1.99 and 3.89 were observed for the linear and cubic spline interpolators, which coincide well with theoretical orders of accuracy. A comparative value for ACGI2 cannot be provided since the concept of motion estimation is not applicable in 1D.

The results of this evaluation can also be verified qualitatively. Figure 4 displays cross-sections from one set of parabolic flow reconstructions. The ACGI2 reconstruction is most similar to the original flow. Cubic splines perform well on the interior, but suffer toward the outskirts as the white contours indicate. By appropriately using a broader region of support, ACGI2 performs accurately on the interior, like cubic splines, while maintaining the advantages that directional interpolation offers near the edges.

Much of the error for the cubic spline interpolation lies outside of the fluid domain, which is better preserved by ACGI2. Considering this type of error is not unrealistic as it is common experimental practice to include boundaries in measurement planes. Nevertheless, when interpolation error was summed only over the known fluid domain, the average MSE for cubic spline interpolation was still larger

than that for ACGI2 by a 19.6% margin. Considering the entire volume this value increased to 58.4%.

3.2 Validation 2

In the second evaluation, three-component velocity fields reconstructed from PIV data with different interpolators were compared to CFD results for flow through a bileaflet mechanical prosthesis. Steady flow through the valve within an in vitro flow chamber (Fig. 5) was sampled with a LaVision stereo PIV system (<http://www.lavision.de>, Lavision, Ypsilanti, MI). Planes were acquired at 3 mm intervals with an in-plane resolution of 0.24 mm for Reynolds numbers ranging from 750 to 4,000. A sodium iodide blood analog solution with a viscosity of 3.5 cP and a density of 1.1 g/cm³ was used to simulate human blood. It was seeded with Plyolite particles to facilitate PIV data acquisition. Chamber and fluid refractive indices were matched. For the CFD component a computational model of the apparatus was generated in SolidWorks based on manufacturer specifications and measurements (<http://www.solidworks.com>, SolidWorks, Concord, MA). Boundary conditions, fluid properties, and Reynolds numbers were specified so as to match the in vitro experimentation. Unsteady simulations using a grid of 1.5×10^6 nodes were carried out using an in-house CFD flow solver, which was previously validated for more than 2000 time steps by Ge et al. (2003), (2005). These unsteady solutions were time averaged to obtain the results that were then used to benchmark interpolation methods.

PIV images extracted from the in vitro model were used to reconstruct the entire fluid domain downstream of the valve for a Reynolds number of 1,200. Both ACGI2 and cubic splines, the most competitive techniques from the first validation, were used to interpolate. Normalized velocity magnitudes from both reconstructions were compared to corresponding values from CFD on a point-by-point basis to determine MSE. Only values within the fluid domain were considered. Divergence values were evaluated on the same point-by-point basis, but compared to the theoretical standard of zero, which is imposed, or nearly so, in CFD. The results of both MSE analyses are given in Table 1.

Fig. 5 Images of **a** the flow chamber and **b** the enclosed bileaflet mechanical valve. In vitro flow through **b** was sampled with PIV and reconstructed with different interpolators for comparison to CFD




Table 1 Comparison between interpolated PIV data and CFD for flow through a mechanical prosthesis

Interpolator	Interpolation MSE	Divergence MSE (1/s ²)
Cubic spline	1.39E-02	61.2
ACGI2	1.02E-02	48.2

The interpolation MSE values correspond to average velocity differentials of 9.87% between CFD and ACGI2, and 12.1% between CFD and cubic splines. While these interpolation errors may seem relatively large, it is noteworthy that the underlying PIV data did not match the corresponding CFD planes exactly. However, under the assumption of strong similarity between the data sources, ACGI2 reconstructed the fluid domain with higher quality based on interpolation MSE. More specifically, interpolation MSE for cubic spline interpolation was 36.3% greater than that for ACGI2, which is an improvement over the 19.6% margin observed in the first validation when errors were summed only over the fluid domain. This improvement is not surprising in the context of a more challenging test case. The object-based foundation of ACGI2 makes it better suited to address complex fluid dynamics than scene-based methods like cubic splines.

The divergence errors for this validation case may seem large as well, but can also be attributed in part to the underlying data as the images were acquired somewhat sparsely (3 mm) and registered manually, and because PIV itself is not divergence-free. Regardless, the divergence errors for cubic spline interpolation were 27.0% greater than those for ACGI2. As a point of reference, the divergence MSE for the CFD solution was 5.65 1/s², within an order of magnitude of the ACGI2 value. Divergence MSE was 8.5 times the CFD value for ACGI2, and 10.9 times that value for cubic splines.

The high-quality performance of ACGI2 can also be observed qualitatively in this experiment. Figure 6 shows velocity magnitude contours in a cutaway view of one jet a short distance downstream of the valve. More significant ledging artifacts are apparent in the cubic spline interpolation as opposed to the ACGI2 version. Here the new method does a better job of preserving the boundaries of the jet fluid structure. Figure 7 displays divergence contour plots for part of a slice near the center plane from both

Fig. 6 Cutaway views of a jet downstream from the bileaflet valve. Velocity magnitudes within the light blue chamber have been normalized by the maximum value in each fluid domain. Shown are **a** the complete CFD solution, **b** the region of interest from CFD, and **c** and **d** that same region from the ACGI2 and cubic spline reconstructions, respectively

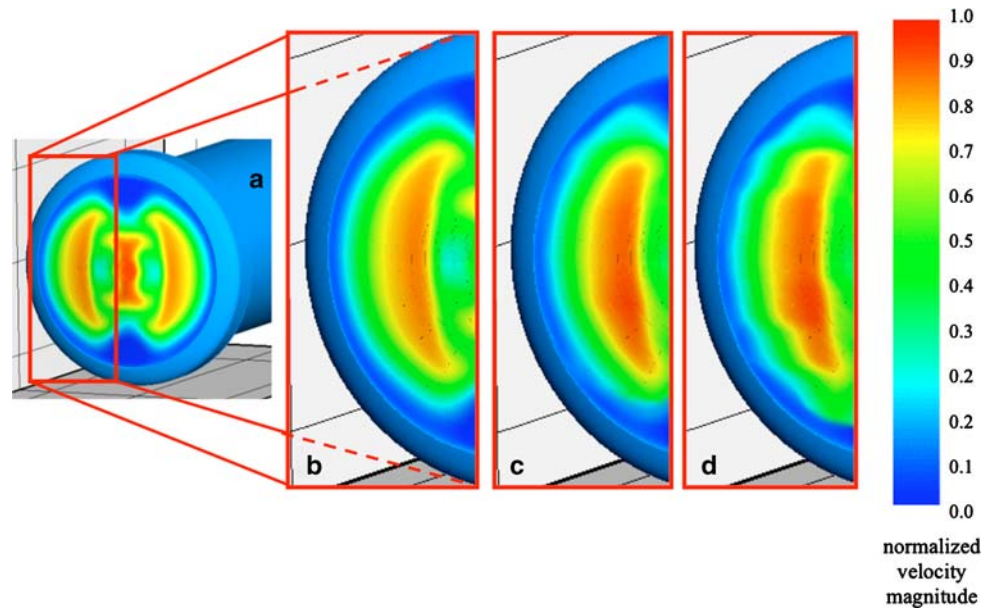



Fig. 7 Color-coded divergence error magnitude plots for part of an interpolated plane near the center plane from **a** the ACGI2 reconstruction and **b** the cubic spline reconstruction



reconstructions. This slice was interpolated at a location equidistant from the center plane and one of its neighbors. Larger divergence errors for the cubic spline interpolation in both the high magnitude jet region and the nearby recirculation zone are apparent.

3.3 Validation 3

Validation 3 is a primarily qualitative exercise that explores the implications of using different interpolators for numerical particle tracking. A surgically modified vascular anatomy based on MR data was fabricated via rapid prototyping for PIV experimentation. Details on data acquisition, reconstruction, and the rapid prototyping process can be found in (Frakes et al. 2008) and (de Zelicourt

et al. 2005). Boundary conditions were taken from PCMR data acquired with these parameters: TR 50 ms, TE 3.8 ms, FOV 250 mm, resolution 0.98 mm, and VENC 80 cm/s. A sodium iodide blood analog solution with a viscosity of 3.5 cP and a density of 1.1 g/cm³ was seeded with red fluorescent particles and circulated through the model under steady flow conditions. As is standard protocol for this type of optical experimentation, refractive indices for the model and fluid were matched (de Zelicourt et al. 2005).

All three velocity components were sampled from the model using a LaVision stereo PIV system at 1 mm intervals with an in-plane resolution of 0.28 mm. Images were reconstructed to form isotropic data volumes with both ACGI2 and cubic splines. In the context of numerical

particle tracking, significant differences were observed. Figure 8a, b display anterior coronal and left sagittal perspectives of the paths followed by 50 particles, tracked back through circulation from the left pulmonary artery (LPA). Paths are red for ACGI2 reconstruction, and blue for cubic splines. The final destinations of the two sets of particles in the LPA are identical, but when paths are integrated back through the respective velocity fields for each, the origins are very different. ACGI2 interpolation shows a majority of the particles originating from the inferior vena cava (IVC), while the cubic spline interpolation shows more particles originating from the superior vena cava (SVC). This fundamental difference is highly relevant from a clinical perspective since even flow distribution to the two outlets is a primary surgical goal (Pekkan et al. 2005).

Ground truth is not available for this in vivo case, so neither solution can be considered more correct. However, divergence errors were 55.6% greater for cubic spline interpolation as compared to ACGI2. This margin is greater than the 27.0% discrepancy observed for the driven cavity problem, reinforcing that ACGI2 performs comparatively better for more geometrically challenging domains with arbitrary shapes.

Although PIV data is not as commonly used for numerical particle tracking as CFD data, many biofluids applications do rely on PIV for this purpose (Raz et al. 2007; Yu and Zhao 1999; Nguyen et al. 2004). It is noteworthy that while the velocity fields in this example were reconstructed with cubic splines in one case and ACGI2 in the other, the paths through these fields were integrated using a separate second-order accurate commercial package (<http://www.tecplot.com>, Tecplot, Bellevue, WA). ACGI2 can be used to integrate particle paths directly, but this methodology was selected in order to put the two competing interpolators on a level playing field. The results




Fig. 8 Anterior coronal (a) and left sagittal (b) perspectives of the paths followed by 50 particles that converge to identical locations in the LPA. Paths for the ACGI2 velocity field reconstruction are indicated in red; those for cubic spline interpolation are represented in blue. Arrows indicate the general direction of flow

demonstrate that the implications of using different interpolators in the context of numerical particle tracking are significant, and especially relevant in this case from a clinical perspective.

This real-world example also highlights the ability of ACGI2 to effectively reconstruct coherent flow. The surgically modified vascular anatomy in Fig. 8 is characterized by highly variable and complex fluid dynamics. Figure 9b displays a vorticity plot for a plane extracted from the model. Plots for interpolated versions of that plane are shown in (c) and (d). The shear layers present in the acquired plane are better preserved with ACGI2 as opposed to cubic spline interpolation. In another test case involving coherent flow, these two interpolators were used to reconstruct planes of velocity data extracted from a CFD solution to the classic driven cavity problem. Cubic splines

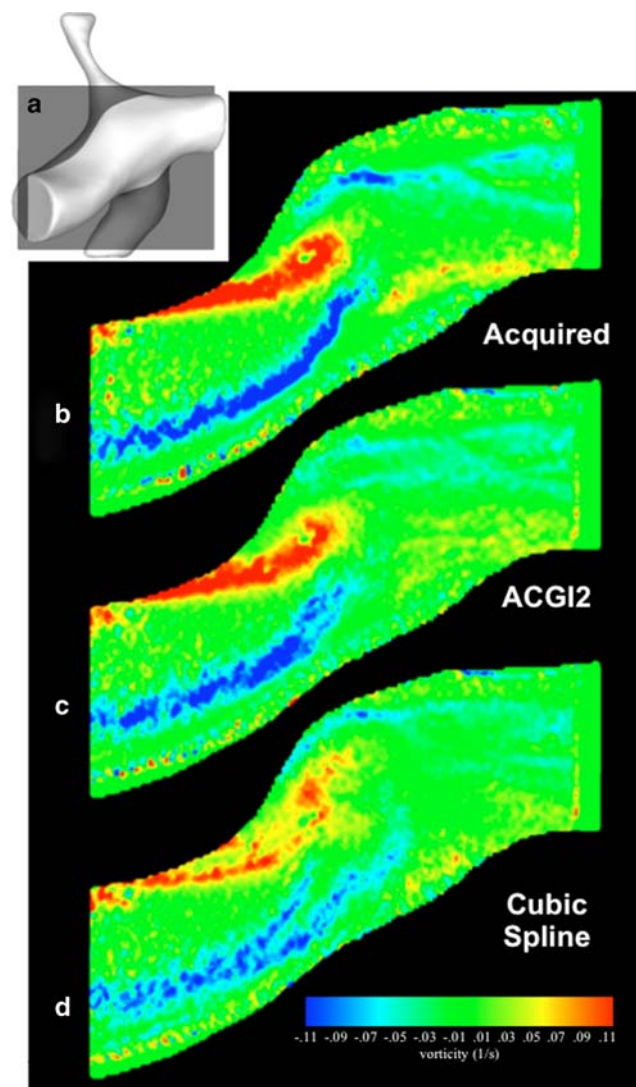


Fig. 9 Vorticity plots for a plane from the anatomical model (a). The plots correspond to the originally acquired data (b), ACGI2 reconstructed data (c), and cubic spline interpolated data (d)

demonstrated 39.7% greater interpolation MSE and 14.7% greater vorticity MSE. High-quality performance in the presence of coherent fluid structures is one feature that makes ACGI2 attractive for 3D experimental applications involving spatially rich flow.

4 Discussion and conclusions

There are several limitations associated with the use of ACGI2 in addressing velocity data interpolation. First, the optical flow equation is based on several assumptions, one of which is that the value at each point within an image remains constant as that point transitions to another image. Clearly this is not the case in the context of fluid flows as velocities change from one measurement plane to the next. However, an explicit solution to the OFCE is not required in the ACGI2 framework; rather the error associated with the OFCE is minimized so that violations of the aforementioned constraint can be accommodated.

Another limitation arises when neighboring images are extremely different and lack similar features to be linked, or when the motion model is unable to detect or express the correct displacements. However, these scenarios are infrequently observed when sampling is performed with sufficient density. Nevertheless, when such circumstances are present a significant decrease in OFCE error cannot be achieved. The result is that a zero value displacement field is used, which translates to a standard cubic spline interpolation. This outcome is achieved through the parameter-based optimization modifications discussed in the Methods section. For these cases ACGI2 can perform no worse than the fallback interpolator, and in the majority of cases it performs better.

The proposed technique has been presented here in the context of reconstructing isotropically resolved data volumes from sparsely acquired measurement planes, because in that role it can add value to a broad range of applications. However, the algorithm is not limited to such applications. ACGI2 addresses interpolation at both discrete and arbitrary coordinates equally well, and can also be applied to problems involving data that are natively isotropic.

The question has been posed, “Why should interpolation between points registered with optical flow provide better quality than interpolation between points registered by coordinates?” When points registered by the latter criterion lie inside and outside of the fluid domain, respectively, the answer is clear. When all points in question are within the fluid domain, as in the vast majority of cases, the answer is less clear but the authors maintain that the same analogy holds. Just as interpolation between points inside and outside of the fluid domain produces inaccurate data,

interpolation between points from unrelated fluid structures produces less accurate data. Accordingly, the registration of similar fluid structures prior to interpolation contributes to improved accuracy, as the results in this paper indicate. For brevity the authors have chosen not to include the results for conventional ACGI here, but it is noteworthy that ACGI2 did outperform that formulation in all of the validations. Future work will explore the incorporation of special divergence-free basis functions into the ACGI2 framework (Lowitzsch 2004).

Current imaging technologies allow the acquisition of high-quality measurement planes from complex experimental flows, but are often limited by lower bounds on plane spacing. Many applications call for more highly resolved 3D fluid velocity information. ACGI2 is a novel solution to this problem that performs with improved quality in the context of fluid flow reconstruction. ACGI2 is built upon the foundation of traditional ACGI but incorporates several modifications that enable it to address velocity interpolation problems more successfully. The linkage of adjacent displacement fields to allow higher-order interpolation and the incorporation of fluids-specific information into the optimization framework contribute to significantly improved results in comparison to the traditional formulation. In multiple theoretical and in vitro validations, ACGI2 outperformed the most popular solution from current literature, cubic spline interpolation, based on interpolation and divergence error metrics. ACGI2 also reconstructed coherent flow structures with higher quality than the competing method. Lastly, ACGI2 performed comparatively better for more challenging cases characterized by complex flow and arbitrary geometries, and for real-world biomedical examples.

Acknowledgments This research was supported by a grant from the National Institutes of Health (BRP grant #R01HL67622). The authors would also like to thank Dr. Mark Fogel of the Children’s Hospital of Philadelphia, Dr. Fotis Sotiropoulos of the University of Minnesota, and Dr. Liang Ge of the Georgia Institute of Technology for their valuable contributions to this work.

References

- Chen MC, Willson AN Jr (2000) Motion-vector optimization of control grid interpolation and overlapped block motion compensation using iterated dynamic programming. *IEEE Trans Imaging Proc* 9:1145–1147
- Derou D, Dinten J, Herault L, Niez J (1995) Physical-model based reconstruction of the global instantaneous velocity field from velocity measurement at a few points. In: *Proceedings: physics-based modeling in computer vision*, p 63
- Elkins CJ, Alley MT (2007) Magnetic resonance velocimetry: applications of magnetic resonance imaging in the measurement of fluid motion. *Exp Fluids* 43:823–858
- Frakes DH, Monaco J, Smith M (2001) Suppression of atmospheric turbulence in video using an adaptive control grid interpolation

- approach. In: Proceedings: international conference on acoustics, speech, and signal processing, May 2001
- Frakes DH, Conrad C, Healy T, Monaco JW, Smith MJT, Fogel M, Sharma S, Yoganathan AP (2003) Application of an adaptive control grid interpolation technique to morphological vascular reconstruction. *IEEE Trans Biomed Eng* 50:197–206
- Frakes DH, Smith MJT, de Zelicourt D, Pekkan K, Yoganathan AP (2004) Three-dimensional velocity field reconstruction. *J Biomech Eng* 126:727–735
- Frakes DH, Pekkan K, Dasi LP, Kitajima HD, Yoganathan AP, Smith MJT (2008) A new adaptive method for registration-based medical image interpolation. *IEEE Trans Med Imaging* 27:370–377
- Ge L, Jones SC, Sotiropoulos F, Healy TM, Yoganathan AP (2003) Numerical simulation of flow in mechanical heart valves: grid resolution and the assumption of flow symmetry. *J Biomech Eng* 125:709–718
- Ge L, Leo HL, Sotiropoulos F, Yoganathan AP (2005) Flow in a mechanical bileaflet heart valve at laminar and near-peak systole flow rates: CFD simulations and experiments. *J Biomech Eng* 127:782–797
- Grevera GJ, Udupa JK (1996) Shape-based interpolation of multi-dimensional grey-level images. *IEEE Trans Med Imaging* 15:881–892
- Horn B (1968) Robot vision. The MIT Press, Cambridge, pp 280–292
- Karatekin O, Wang FY, Charbonnier JM (1998) Characterization of a bluff body wake using LDV and PIV techniques. In: Ninth international symposium on applications of laser techniques to fluid mechanics, July 1998
- Kimura Y, Herring JR (1996) Diffusion in stably stratified turbulence. *J Fluid Mech* 328:253–269
- Lowitzsch S (2004) Approximation and interpolation employing divergence-free radial basis functions with applications. Ph.D. thesis, Texas A&M University
- Melnikov DP, Shevtsova VM (2005) Liquid particles tracing in three-dimensional buoyancy driven flows. *Fluid Dyn Mater Process* 1:189–199
- Monaco JW (1997) Motion models for video applications. Ph.D. thesis, Georgia Institute of Technology
- Nguyen TT, Biadillah Y, Mongrain R, Brunette J, Tardif JC, Bertrand OF (2004) A method for matching the refractive index and kinematic viscosity of a blood analog for flow visualization in hydraulic cardiovascular models. *J Biomech Eng* 126:529–535
- Ozturk C, Derbyshire JA, McVeigh ERM (2003) Estimating motion from MRI data. *Proc IEEE* 91:1627–1648
- Pekkan K, Kitajima HD, de Zelicourt D, Forbess JM, Parks WJ, Fogel MA, Sharma S, Kanter KR, Frakes DH, Yoganathan AP (2005) Total cavopulmonary connection flow with functional left pulmonary artery stenosis: angioplasty and fenestration in vitro. *Circulation* 112:3264–3271
- Penney GP, Schnabel JA, Rueckert D, Viergever MA, Niessen WJ (2004) Registration-based interpolation. *IEEE Trans Med Imaging* 23:922–926
- Raz S, Einav S, Alemu Y, Bluestein D (2007) DPIV prediction of flow induced platelet activation—comparison to numerical predictions. *Ann Biomed Eng* 35:493–504
- Sakakibara J, Nakagawa M, Yoshida M (2004) Stereo-PIV study of flow around a maneuvering fish. *Exp Fluids* 36:282–293
- Yeung PK (2002) Lagrangian investigations of turbulence. *Annu Rev Fluid Mech* 34:115–142
- Yeung PK, Pope SB (1988) An algorithm for tracking fluid particles in numerical simulations of homogeneous turbulence. *J Comput Phys* 79:373–416
- Yu SCM, Zhao JB (1999) Steady and pulsating flow characteristics in straight tubes with and without a lateral circular protrusion. *Exp Fluids* 26:505–512
- de Zelicourt D, Pekkan K, Kitajima HD, Frakes DH, Yoganathan AP (2005) Single-step stereolithography of complex anatomical models for optical flow measurements. *J Biomech Eng* 127:205–207
- Zhao W, Sawhney H (2002) Is super-resolution with optical flow feasible? In: Proceedings: European conference on computer vision, vol 1