

# High-Speed Registration of Three- and Four-dimensional Medical Images by Using Voxel Similarity<sup>1</sup>

Raj Shekhar, PhD • Vladimir Zagrodsky, PhD • Carlos R. Castro-Pareja, MS • Vivek Walimbe, BE • Jogikal M. Jagadeesh, PhD

A generalized, accurate, automatic, retrospective method of image registration for three-dimensional images has been developed. The method is based on mutual information, a specific measure of voxel similarity, and is applicable to a wide range of imaging modalities and organs, rigid or deformable. A drawback of mutual information-based image registration is long execution times. To overcome the speed problem, low-cost, customized hardware to accelerate this computationally intensive task was developed. Individual hardware accelerator units (each, in principle, 25-fold faster than a comparable software implementation) can be concatenated to perform image registration at any user-desired speed. A first-generation prototype board with two processing units provided a 12- to 16-fold increase in speed. Enhancements for increasing the speed further are being developed. These advances have enabled many nontraditional applications of image registration and have made the traditional applications more efficient. Clinical applications include fusion of computed tomographic (CT), magnetic resonance, and positron emission tomographic (PET) images of the brain; fusion of whole-body CT and PET images; fusion of four-dimensional spatiotemporal ultrasonographic (US) and single photon emission CT images of the heart; and correction of misalignment between pre- and poststress four-dimensional US images.

©RSNA, 2003

**Abbreviations:** FAIR = fast automatic image registration, MI = mutual information, 2D = two-dimensional, 3D = three-dimensional, 4D = four-dimensional

**Index terms:** Data fusion • Dual-modality imaging • Images, fusion

**RadioGraphics 2003;** 23:1673–1681 • **Published online** 10.1148/rg.236035041

<sup>1</sup>From the Department of Biomedical Engineering/ND20, Lerner Research Institute, Cleveland Clinic Foundation, 9500 Euclid Ave, Cleveland, OH 44195 (R.S., V.Z., C.R.C.P., V.W.); and the Department of Electrical Engineering (C.R.C.P., J.M.J.) and Biomedical Engineering Center (V.W.), Ohio State University, Columbus. Presented as an *infoRAD* exhibit at the 2002 RSNA scientific assembly. Received February 21, 2003; revision requested May 2 and received June 12; accepted June 23. Supported in part by research grant DAMD17-99-1-9034 from the Department of Defense and grant RG-01-0071 from the Whitaker Foundation. **Address correspondence to** R.S. (e-mail: [shekhar@bme.ri.ccf.org](mailto:shekhar@bme.ri.ccf.org)).

©RSNA, 2003

## Introduction

Image registration involves aligning two or more images before they can be meaningfully overlaid or compared. Medical image registration is applied to three major areas: multimodality fusion (overlying) of images that provide complementary information (typically structural and functional); serial comparison to quantify disease progression or regression, especially in response to treatment; and intersubject registration, which is aimed at creating a normalized atlas representative of a patient population (1,2).

The focus of the work presented in this article is the registration of three-dimensional (3D) and four-dimensional (4D) images. Whereas static organs can be represented adequately by using 3D images, 4D (3D space plus time) images are essential for imaging dynamic organs such as the heart. In this section, we review the suitability of existing image registration methods for a range of 3D and 4D medical images and contend that the voxel similarity–based approach is the most general and powerful approach.

Image registration has most often been applied to neuroimaging (2,3). Since the shape and size of the brain stay fixed because of the bony skull surrounding it, simple rigid-body registration involving only six transformation parameters (three translations and three rotations) suffices. This registration is typically carried out by matching at least four homologous point landmarks in the images, where the point landmarks are external markers screwed into the skull or built into a stereotactic frame fitted to the patient. Its prospective nature (the burden of placing markers) makes the clinical protocol lengthy and complicated. Because it is invasive, the placement of markers is not allowed for diagnostic purposes at most hospitals. Registration is thus limited to neurosurgical cases. Even in neurosurgical patients, possibly better care arising from fusing positron emission tomographic (PET) images with computed tomographic (CT) and magnetic resonance (MR) images is not feasible because the frame and the markers do not show in PET images. Unlike CT and MR imaging, PET allows definitive identification of viable tumors. Clearly, marker-based registration has limited usefulness in multimodality fusion.

Diagnosis and treatment of diseases in most organs, not just the brain, stand to benefit from image registration. Most other organs are also deformable or nonrigid. Prospective marker-based registration is both inaccurate and cumbersome for deformable organs because the organs can move relative to externally placed markers and because many more markers are needed to solve for a large number of parameters associated with non-rigid-body registration. Both rigid-body and non-rigid-body image registration have been attempted, retrospectively, by matching internal image features (points, contours, and surfaces) identified by means of segmentation (4). However, reliable automatic segmentation is extraordinarily complex. Consequently, segmentation generally involves manual steps, and segmentation-based registration cannot be fully automatic. The accuracy of segmentation always remains suspect, and additional errors may be introduced because the location of a specific (segmented) feature can vary both within and between modalities.

The relatively new voxel similarity–based approach provides the most flexible and general theoretical framework for most registration tasks (2). Its accuracy over the surface segmentation–based approach for multimodality registration of brain images has been shown (5). There are also no extrinsic limits on the accuracy and speed and no theoretical limit on the nature of the transformation (rigid or nonrigid) involved. Further, a voxel similarity–based technique can be made fully automatic. Various measures of voxel similarity have been compared, and mutual information (MI) (6,7) has emerged as the most accurate and robust measure (8,9). Successful application of MI-based registration to brain (7,8,10), liver (9), breast (11), chest (12), and cardiac (13–15) images from various modalities has been reported. MI-based registration is also featured in the Insight Segmentation and Registration Toolkit (16), a large-scale, multicenter, open-source software system initiative undertaken by the National Library of Medicine of the National Institutes of Health.

A drawback of MI-based registration is its long execution time. MI is essentially a cost function, which is maximized iteratively by using a standard optimization algorithm. Repeated MI computation, which is memory access–intensive, does not benefit from the cache-based memory archi-

tures present in modern computers. Since memory access time does not evolve according to Moore's law, MI-based registration is not expected to speed up significantly in the near future because of enhanced computer architectures. Prompted by these observations, we have developed an optimized hardware, the fast automatic image registration (FAIR) architecture, to speed up MI-based registration.

In this article, we describe the basics of MI-based image registration followed by an explanation of what limits the performance of the algorithm and how the FAIR architecture overcomes this performance bottleneck. We then describe the application of our high-speed image registration technology to four specific clinical problems. We conclude with discussions on current results and future directions.

### MI-based Image Registration

MI-based image registration, as introduced by Wells et al (6) and Maes et al (7), attempts to maximize the MI between a reference image ( $A$ ) and a floating image ( $B$ ). During the iterative process, the reference image is kept stationary, whereas candidate transformations are applied to the floating image. The registration is said to be completed when the MI is maximized. The corresponding transformation  $\hat{T}$  can then be written as follows:

$$\hat{T} = \arg \max_T MI[A, T(B)]. \quad (1)$$

The MI is calculated from individual and joint entropies as follows:

$$MI(A, B) = H(A) + H(B) - H(A, B). \quad (2)$$

The individual entropies,  $H(A)$  and  $H(B)$ , and the joint entropy,  $H(A, B)$ , in turn are calculated as follows:

$$H(A) = - \sum_a p_A(a) \ln p_A(a), \quad (3)$$

$$H(B) = - \sum_b p_B(b) \ln p_B(b), \quad (4)$$

and

$$H(A, B) = - \sum_{a,b} p_{AB}(a, b) \ln p_{AB}(a, b), \quad (5)$$

where  $p_A(a)$  and  $p_B(b)$  represent individual probability density functions and  $p_{AB}(a, b)$  represents the joint probability density function. The normalized mutual histogram provides a practical estimate of the joint probability density function. Individual probability functions,  $p_A(a)$  and  $p_B(b)$ , are then determined by summing  $p_{AB}(a, b)$  along the  $b$  and  $a$  axes, respectively. Physically, the goal in MI-based image registration is to minimize the spread of values within the mutual histogram, which in turn minimizes the joint entropy and maximizes the MI (see Eq [2]).

### Numerical Complexity and the FAIR Architecture

MI-based registration usually requires hundreds of iterations (MI evaluations), depending on the image content, the transformation complexity, the degree of initial misalignment, and the optimization algorithm used to maximize the MI function. To analyze the performance bottleneck, we treat the MI calculation as being composed of two steps: (a) computation of the mutual histogram and (b) computation of the joint and individual entropies from the mutual histogram leading to calculation of the MI (Eq [3–5] followed by Eq [2]). Note that the first step involves accessing 3D images (arrays), whereas the second step involves processing the mutual histogram, a two-dimensional (2D) array. With current microprocessor speeds on the order of 1 GHz, the overall MI calculation time depends almost exclusively on the memory access-intensive first step.

Computation of the mutual histogram involves finding a matching voxel in the reference image for each floating image voxel after application of the candidate transformation. Since a transformed floating image voxel may not necessarily coincide with a reference image voxel, eight neighboring voxels are interpolated. We use partial-volume interpolation for its well-documented stability (reliable convergence) and performance (subvoxel accuracy) in MI-based registration (7). When partial-volume interpolation is used, the mutual histogram is built up by incrementing the bins corresponding to the eight reference image voxel values by their fractional contributions

**Figure 1.** Pseudocode of the mutual histogram (*MH*) calculation and the associated memory accesses. *PV* = partial volume.

for each voxel in B	
read its intensity	<b>(1 memory access)</b>
find its transformed location in A	
if transformed location is inside A	
read intensities of 8 neighbors	<b>(8 memory accesses)</b>
calculate 8 PV interpolation weights	
read 8 MH entries	<b>(8 memory accesses)</b>
increment those entries	
write 8 updated MH entries	<b>(8 memory accesses)</b>

(weights) for each floating image voxel. Floating image voxels landing outside the bounds of the reference image are ignored. Figure 1 shows the pseudocode of this algorithm and the number of memory accesses against each step.

In a standard (single-processor) software implementation of the algorithm, the described steps must be completed before the next voxel can be processed. The FAIR architecture, which has been described in detail elsewhere (17), speeds up computation of the mutual histogram by use of pipelining, parallel memory access, and distributed computing. The combination of pipelining and parallel memory access effectively reduces the time of 25 memory accesses to that of one memory access. Therefore, computation of the mutual histogram, in theory, is accelerated by a factor of 25. Further increase in the speed can be gained by distributed computing, which divides the computational burden equally among a number of units, each processing a fraction of the floating image.

A prototype board with two processing units was fabricated and used for demonstration. The hardware implementation was 12–16 times faster than an equivalent software implementation on a dual 1-GHz Pentium III (Intel, Santa Clara, Calif) computer with 2 Gbytes of memory, depending on the image size (Table 1). Although no practical implementation can achieve the theoretically predicted maximum speed, various practical issues, as explained in the Discussion section, prevented us from achieving a higher speed. These issues are being addressed in our second-generation prototype board.

**Table 1**  
**Typical MI Calculation Times in Software and Hardware and the Acceleration Ratio**

Image Size (voxels)	MI Calculation Times (msec)		Acceleration Ratio
	In Software	In Hardware	
64 <sup>3</sup>	430	36	11.9
128 <sup>3</sup>	3,500	220	15.9
256 <sup>3</sup>	28,000	1,700	16.5


## Results

We present four clinical applications whose success is critically dependent on image registration. A detailed description and validation of these studies, which are still ongoing, is not intended here. Instead, these applications are presented to convey the enabling nature of our high-speed MI-based image registration technology. Without automatic and fast image registration, these applications may not be suitable for clinical use.

### Fusion of CT, MR, and PET Images of the Brain

Registration and fusion of only the MR and CT image pair is performed routinely in planning radiation treatment of brain tumors. However, MR and CT images may be inconclusive in differentiation of benign masses or necrotic tissue from active tumors. An additional PET scan can help resolve the issue and aid in preventing unnecessary treatment, but the existing frame-based registration method does not extend to PET images. Hardware-accelerated MI-based registration enables this application by first allowing CT-PET





**Figure 2.** Registration and fusion of CT (a) and PET (b) images of the brain to create a multimodality overlay image (c). The PET image (b) shows a tumor (within white box); the CT image (a) shows its location within the brain.

**Table 2**  
**Comparison of Software and Hardware Speeds for Registration of Brain Images**


Image Registration Task	Execution Time of Software (min)	Execution Time of Hardware (sec)
CT-MR	14–18	60–80
CT-PET	11–15	45–60
MR-PET	18–22	80–100

and MR-PET image registration, which cannot be performed currently. Second, the short execution time of the hardware implementation, as we report later, makes it feasible to incorporate these registration tasks into the current clinical protocol efficiently and without adding slow and tedious manual steps.

MI-based registration works for any combination of MR, CT, and PET images without requiring that the patient wear a frame. Figure 2 shows an example of CT and PET image fusion performed with our independent implementation of the MI-based image registration algorithm. The functional PET image allows confirmation of a recurrent active tumor in the left lateral occipital cortex and the need for therapy; the fused image helps localize the tumor precisely. The times for performing CT-MR, CT-PET, and MR-PET

image registration by using both the software and the hardware implementations are reported in Table 2. The hardware solution accelerated MI-based image registration by a factor of 12–16. The exact length of the execution depended on the starting misalignment and the image size (number of voxels). The typical size of the CT, MR, and PET images used was  $256 \times 256 \times 62$ – $80$ ,  $256 \times 256 \times 126$ – $154$ , and  $128 \times 128 \times 63$  voxels, respectively.

The accuracy of use of MI for brain CT-MR and MR-PET image registration has already been demonstrated (7,8,10). In an independent validation study that involved eight patients undergoing gamma knife surgery, we found the mean registration error to be 2.4 mm for CT-MR, 4.0 mm for MR-PET, and 3.3 mm for CT-PET image registration. Subvoxel accuracy, or a residual error less than the dimensions of the PET voxel (approximately 6 mm), was achieved for registrations that involved PET. The accuracy of CT-MR image registration may be improved by correcting MR images for spatial distortions (10). Overall, the initial indications of our ongoing validation study are that our implementation of MI-based registration is comparable in accuracy to external marker-based registration, the perceived standard of reference.



**Figure 3.** Coronal (a) and sagittal (b) fused CT-PET images of the torso for identification and localization of tumors.

### Fusion of Whole-Body CT and PET Images

The clinical motivation for fusing whole-body CT and PET images is identical to that for the brain application described earlier. For treatment of thoracic and abdominal tumors, CT and PET images are routinely acquired and then interpreted independently. Registering and fusing the two modalities (Fig 3) helps localize the active tumors that PET shows with respect to the anatomy as seen at CT. However, the deformable nature of the torso makes the registration task more challenging compared to registration of brain images. This is also a task that cannot be performed accurately by using external markers, whereas use of internal anatomic landmarks makes it excessively labor-intensive and cumbersome for routine clinical use.

However, it is possible to register such images accurately and conveniently by using the MI-based approach. Furthermore, our high-speed implementation can provide the needed clinical efficiency. When a nonrigid transformation model that included scaling was used, the time of whole-body CT-PET image registration in software was 11–17 minutes; the hardware implementation


reduced this time to 45–70 seconds. The typical size of the CT and PET images was  $256 \times 256 \times 65$ – $100$  and  $128 \times 128 \times 135$ – $165$  voxels, respectively. Although our independent validation of this application is still ongoing, Mattes et al (12) reported achieving accurate registration for this combination of chest images by using the MI-based algorithm.

### Fusion of Cardiac SPECT and US Images


Single photon emission CT (SPECT) shows myocardial perfusion and has high sensitivity. Ultrasonography (US) shows myocardial wall thickness and motion and is highly specific. Fusing the myocardial perfusion data from SPECT with the wall structure and motion information from US could improve the sensitivity and specificity of diagnosing myocardial wall disease, which could then help clinicians make better revascularization decisions and reduce unnecessary catheterizations. An example of such fusion is shown in Figure 4. Because the heart moves and has few well-defined anatomic landmarks, use of traditional registration methods is difficult and inaccurate. In a preliminary investigation, we have shown that use of MI is an effective and attractive alternative for this image registration task (18).

This and similar cardiac applications require spatial registration of 10–30 temporally matched





**Figure 4.** Registration and fusion of real-time 3D US (a) and gated SPECT (b) images of the heart to create a multimodality overlay image (c). The images show a long-axis view of the heart in diastole, as extracted from the 4D images.



**Figure 5.** Correction of misalignment between prestress (a) and poststress (b, c) 3D US images of the heart in systole. Despite stress-induced size differences, the left ventricle (yellow outline) is better aligned on the poststress image obtained with registration (c) than on the one obtained without registration (b).

3D image pairs per patient, making high-speed image registration crucial to their success. In four cases, the time to register a single 3D US-SPECT frame pair ranged between 7 and 10 minutes in software. The size of a 3D US image was  $128 \times 128 \times 128$  voxels, and that of a SPECT frame was  $64 \times 64 \times 27$ – $42$  voxels. Because 16–25 image pairs needed to be registered, the total time for 4D image registration often exceeded 2 hours. Although this is an offline diagnostic application, such long registration times are clinically unacceptable. Our hardware implementation, which reduced the registration time to 24–35 seconds per frame pair, is key to clinical implementation of this application.

#### Alignment Correction of Pre- and Poststress 3D Echocardiograms

Diagnosing myocardial ischemia accurately is important because it may save many lives by allowing narrowed coronary vessels to be reopened

before the myocardium becomes irreversibly scarred. Myocardial ischemia is commonly diagnosed with stress echocardiography, in which the response of the heart to exercise or other forms of stress is measured by comparing the active and resting phases of the left ventricular wall on a series of US images. A common problem in stress echocardiography is that the rest and stress image pairs are misaligned. We use real-time 3D US, instead of conventional 2D US, in conjunction with the existing stress protocol and use MI-based registration to correct for misalignment between pre- and poststress images due to probe placement errors while preserving stress-induced changes. An example of such alignment correction is shown in Figure 5. As 3D US becomes prevalent, this emerging 3D stress echocardiography approach can conceivably replace the existing procedure.

We have demonstrated that MI-based image registration is both possible and accurate for intramodality registration of 3D US images (13). The time to nonrigidly register a pair of pre- and poststress images, each sized  $128 \times 128 \times 128$  voxels, was 7–10 minutes and 25–35 seconds with the software and hardware implementations, respectively. Furthermore, 12–15 frame pairs spanning a complete cardiac cycle needed to be registered in this application. As with the previous cardiac application, the current 3D stress echocardiography application is enabled by automatic MI-based image registration, whose efficient, high-speed implementation is equally critical for clinical deployment.

### Discussion

The benefits of image registration in multimodality fusion and in single-modality serial comparison are well demonstrated and validated methods exist (2,10,15), but the clinical acceptance of image registration has been low. Deficiencies in the existing methods, which are typically slow, specific, and tedious, are responsible for the low acceptance. Furthermore, existing methods prolong existing protocols by requiring extra steps such as placing external markers. As reported in the literature, MI-based image registration is an attractive alternative to existing methods because it is versatile, accurate, and automatic and, as shown herein, can be performed quickly.

Successful application of MI-based registration to images of the brain (7,8,10), liver (9), breast (11), chest (12), and heart (13–15) with accurate results has been reported. Furthermore, the method has been shown to be applicable to a wide variety of registration tasks. The versatility of MI-based image registration can also be concluded from the four model applications described herein. The same method worked effectively to register images of different regions (the brain, whole body, and heart) and different modalities (CT, MR, PET, SPECT, and US). It also handled rigid and deformable organs well. Moreover, the method is general enough to apply to multimodality as well as single-modality images.

MI-based registration is both accurate and automatic. Rigid-body registration between a structural (MR or CT) and a functional (PET or

SPECT) image has been validated and widely reported. In most cases, subvoxel accuracy (accuracy surpassing the voxel dimensions) was achieved (5,8,10). We have performed extensive validation of US-US image registration and have achieved subvoxel accuracy despite the low spatial resolution of real-time 3D cardiac US images (13). We have also extensively investigated the capture range, the largest starting misalignment from which MI-based registration succeeds. Fortunately, the starting misalignment in most applications is often well within the capture range. If not, a priori knowledge combined with fast, automatic methods (19) can be used to bring the two volumes into approximate alignment. We thus conclude that MI-based registration can be fully automatic.

Speed is the only apparent drawback of MI-based image registration. However, using our FAIR architecture, we showed that the algorithm can be accelerated to operate at a higher speed. By exploiting the scalability of our hardware solution, achievement of any user-defined speed is possible by concatenating multiple processing units. Automated high-speed image registration will improve efficiency and provide ease of use by eliminating complicated (and invasive) imaging protocols and labor-intensive steps such as identifying and matching markers or other image features. Speed should also improve the applicability of the method to emerging 4D (3D + time) cardiac applications. In cardiac applications, one often needs to register 10–30 image pairs corresponding to various phases of the cardiac cycle. Speed is essential to finish the registration task in a reasonable time.

In our current hardware implementation, the achieved speed increase was approximately three times smaller than what was theoretically expected. Various fabrication factors including hardware noise, which forced us to operate at a lower clock frequency than intended, contributed to this performance gap. These fabrication problems are being addressed in the next-generation prototype boards, which are expected to be approximately two orders of magnitude faster.

Although we have presented only four applications, automatic high-speed image registration has endless possibilities. A possible application is in surgical navigation, wherein intraoperative 2D or 3D US images may be registered with preoperative CT or MR images in real time to update a



preoperative plan “on the fly.” Likewise, in patients receiving fractionated radiation treatment, registering preoperative CT images with a low-dose CT or 3D US image obtained per patient visit may be used to direct the treatment precisely to the intended target. Yet another application is real-time display of cardiac US images overlaid on previously acquired gated SPECT images. Possible applications of high-speed 3D image registration are limited only by the human imagination.

## Conclusions

Voxel similarity-based image registration has been gaining popularity. The specific measure of voxel similarity that has attracted the most attention is MI. A growing body of literature confirms that use of MI is the most general, accurate, reliable, and fully automatic approach to image registration available currently. However, MI-based 3D image registration is computationally intensive, and long execution time may well impede the growth and acceptance of this promising method. To overcome the speed problem, we have developed a novel hardware accelerator, the FAIR architecture. We have presented the application of our high-speed method to four diverse clinical problems. We believe that versatility, accuracy, speed, and automation are keys to creating effective image registration-based applications and increasing their acceptance by the clinical community.

**Acknowledgments:** The authors thank the Departments of Cardiovascular Medicine and Radiation Oncology and the Division of Molecular and Functional Imaging within the Department of Radiology at the Cleveland Clinic Foundation for providing the images used in the present work.

## References

1. Bankman IN. Handbook of medical imaging: processing and analysis. San Diego, Calif: Academic Press, 2000.
2. Maintz JB, Viergever MA. A survey of medical image registration. *Med Image Anal* 1998; 2:1–36.
3. Treves S, Mitchell K, Habboush I. Three dimensional image alignment, registration and fusion. *Q J Nucl Med* 1998; 42:83–92.
4. Pelizzari CA, Chen GT, Spelbring DR, Weichselbaum RR, Chen CT. Accurate three-dimensional registration of CT, PET, and/or MR images of the brain. *J Comput Assist Tomogr* 1989; 13:20–26.
5. West J, Fitzpatrick JM, Wang MY, et al. Retrospective intermodality registration techniques for images of the head: surface-based versus volume-based. *IEEE Trans Med Imaging* 1999; 18:144–150.
6. Wells WM, Viola P, Atsumi H, Nakajima S, Kikinis R. Multi-modal volume registration by maximization of mutual information. *Med Image Anal* 1996; 1:35–51.
7. Maes F, Collignon A, Vandermeulen D, Marchal G, Suetens P. Multimodality image registration by maximization of mutual information. *IEEE Trans Med Imaging* 1997; 16:187–198.
8. Studholme C, Hill DL, Hawkes DJ. Automated three-dimensional registration of magnetic resonance and positron emission tomography brain images by multiresolution optimization of voxel similarity measures. *Med Phys* 1997; 24:25–35.
9. Carrillo A, Duerk JL, Lewin JS, Wilson DL. Semi-automatic 3-D image registration as applied to interventional MRI liver cancer treatment. *IEEE Trans Med Imaging* 2000; 19:175–185.
10. West J, Fitzpatrick JM, Wang MY, et al. Comparison and evaluation of retrospective intermodality brain image registration techniques. *J Comput Assist Tomogr* 1997; 21:554–566.
11. Meyer CR, Boes JL, Kim B, et al. Semiautomatic registration of volumetric ultrasound scans. *Ultrasound Med Biol* 1999; 25:339–347.
12. Mattes D, Haynor DR, Vesselle H, Lewellen TK, Eubank W. PET-CT image registration in the chest using free-form deformations. *IEEE Trans Med Imaging* 2003; 22:120–128.
13. Shekhar R, Zagrodsky V. Mutual information-based rigid and nonrigid registration of ultrasound volumes. *IEEE Trans Med Imaging* 2002; 21:9–22.
14. Shekhar R, Zagrodsky V, Garcia M, Thomas JD. 3D stress echocardiography: a novel application based on registration of real-time 3D ultrasound images. In: *Proceedings of Computer-assisted Radiology and Surgery 2002*. Berlin, Germany: Springer-Verlag, 2002; 873–878.
15. Makela T, Clarysse P, Sipila O, et al. A review of cardiac image registration methods. *IEEE Trans Med Imaging* 2002; 21:1011–1021.
16. NLM Insight Segmentation and Registration Toolkit. Available at: [www.itk.org](http://www.itk.org).
17. Castro-Pareja CR, Jagadeesh JM, Shekhar R. FAIR: a hardware architecture for real-time 3-D image registration. *IEEE Trans Inf Technol Biomed* (in press).
18. Walimbe V, Zagrodsky V, Raja S, et al. Mutual information-based multimodality registration of cardiac ultrasound and SPECT images. *Int J Cardiovasc Imaging* (in press).
19. Dhawan AP, Arata LK, Levy AV, Mantil J. Iterative principal axes registration method for analysis of MR-PET brain images. *IEEE Trans Biomed Eng* 1995; 42:1079–1087.