INSTITUTE OF PHYSICS PUBLISHING PHYSICS IN MEDICINE AND BIOLOGY

Phys. Med. Biol. **50** (2005) 1295–1312 doi:10.1088/0031-9155/50/6/017

A lesion stabilization method for coronary angiography

Normand Robert, Philip T Komljenovic, Ryan Grant, Marshall S Sussman and J A Rowlands

Sunnybrook and Women's Health Science Centre, University of Toronto, Room S632, 2075 Bayview Avenue, Toronto M4N 3M5, Canada

E-mail: robertn@cardioview.com

Received 10 August 2004, in final form 20 January 2005
Published 2 March 2005
Online at stacks.iop.org/PMB/50/1295

Abstract

A method to make a coronary artery segment of interest appear stationary when viewing a sequence of angiographic images is proposed. The purpose of this method is to facilitate the assessment of lesions caused by coronary artery disease by improving detectability. A description of the stabilization algorithm based on template matching is given. Stabilization was performed on 41 clinical coronary angiograms exhibiting various stenoses and was successful in 39/41 cases. A quantitative analysis of stabilization errors was performed by introducing simulated moving vessels of decreasing contrast into sequences of clinical images.

1. Introduction

Coronary artery disease (CAD), which causes thickening of the walls of arteries, accounted for 1 in 5 deaths in the US in 1998 (*Heart and Stroke Statistical Update* 2001). This disease gives rise to lesions where the lumen cross-section is reduced thereby impeding blood flow. Thus assessment of artery geometry using angiography plays an important role in the diagnosis and management of CAD.

Lesions are small features whose extent is often more difficult to assess than adjacent normal vessel segments due to their lower contrast caused by a reduction in lumenal cross-section (Wright *et al* 1985). In addition, some lesions may be ruptured allowing contrast agent to leak into the plaque making the lumen boundary difficult to assess (Levin and Gardiner 1992). Reanalysed lesions can also exhibit multiple channels too small to be distinguished individually. Assessment of low contrast lesions is further impaired by: (1) rapid motion of the arteries, (2) x-ray noise and (3) signals due to *other structures* such as the ribs, diaphragm and spine. We now examine how to reduce these sources of image degradation.

Motion has generally been shown to reduce visual acuity through mechanisms such as sensory inhibition (Sekuler and Blake 1990). In the case of x-ray images displaying small cylinders mimicking vessels, detectability is reduced at velocities greater than 2.9° s^{-1} (Xue and Wilson 1998). Coronary arteries have peak velocities of up to 10 cm s^{-1} (Hofman *et al* 1998, Wang *et al* 1999, Achenbach *et al* 2000) resulting in peak angular velocities of 20° s^{-1} from the perspective of the observer[1]. This problem cannot simply be solved by stopping the playback of an angiographic image sequence to examine a particular frame because the perceived x-ray noise is greatly increased when viewing a single frame (Wilson *et al* 1994). This is because observers can perform temporal integration of visual information for periods of 750 (Eckstein *et al* 1996) to 1500 ms (Whiting *et al* 1991) and this ability results in an improvement in the signal to noise when viewing a sequence of images because x-ray noise is largely statistically independent of frame to frame. One cannot increase the exposure to a particular frame to reduce the noise because one does not know *a priori* which frame will be of greater interest. Increasing the exposure to all frames in case one decides to subsequently pause playback in some yet undetermined frame is dose inefficient and some lengthy catheterization procedures can already cause somatic effects (Rosenthal *et al* 1998, Wexler 1995, Berthelsen and Cederblad 1991).

Observers must contend with viewing coronary arteries which are moving rapidly and therefore harder to see while *other structures* such as the ribs and the diaphragm which are less important move comparatively slowly. Cardiac catheterization is unique in that signals due to these other structures cannot in general be removed using digital subtraction angiography (DSA) because of cardiac and respiratory motions (Meijering *et al* 1999). It would be preferable if an artery segment of interest was stationary and the other structures were moving instead. This behaviour can be achieved by shifting each frame during playback to stabilize the segment of interest. Stabilization of a vessel segment would eliminate the need to stop image playback which makes noise more noticeable. Manual stabilization of a sequence of frames (performed, for example, by identifying the feature of interest in each frame with the aid of a mouse) showing a moving x-ray vessel phantom over a uniform background has been shown to increase observer detectability, reduce the time required for detection and uncertainty in the presence of small lumenal morphological features (Eigler *et al* 1994).

Our goal is to develop a method to automate the stabilization of vessels and thus facilitate its use. Stabilization can be used to examine lesions during diagnostic catheterization, percutaneous trans-lumenal interventions (PCI) and following these procedures, during reporting by the interventionist. We describe an automated method based on template matching to rapidly register a sequence of images such that a feature of interest designated by the observer in a single frame (with the aid of a mouse) appears stationary in the subsequent display of the sequence. A measure of tracking accuracy of the automated method relative to manual motion measurements is computed for a training set of cine runs. This measure is used to optimize the automated method through the introduction of constraints on the interframe motion vectors between successive frames and as well as other modifications described below. Finally, the optimized method is applied on a new set of angiograms to ensure that the optimization is valid on any random samples of cine runs and not only the training set.

In order to gain clinical acceptance, the method must not only be simple and fast but also exhibit a high success rate and failures should be obvious. In other words, the method should not alter diagnostic information in a manner that is imperceptible. We will demonstrate that these goals have been achieved with our approach.

[1] Angular velocity calculation assumes an x-ray detector with a 21 cm field of view, 44 cm display device with the observer 50 cm away.

2. Theory

2.1. Image registration

Lesion stabilization is performed by specifying a region encompassing an artery segment of interest in one of a series of frames thereby defining a template image. The feature is searched in the other frames by a process known as template matching that yields measures of interframe motion. These measures are used to register or position the feature in the same location from frame to frame so that it appears stationary.

Image registration has been used extensively to perform retrospective motion correction of the mask image in DSA for the purpose of reducing subtraction artefacts (Meijering *et al* 1999). Using these methods, a sub-region of an image is used as a template that is compared against a target image. A measure of similarity is computed between the target and template images for different amounts of spatial shift of the template relative to the target. Other transformations such as rotation and scale changes could be applied as well. The transformation yielding the best index of similarity is retained.

During image acquisition, the x-ray intensity may change due to the action of the automatic brightness control and diffusion of the contrast agent may occur. These changes have a multiplicative effect on the image (Fitzpatrick *et al* 1988, Cox and de Jager 1994) and therefore, a registration method that is insensitive to multiplicative changes between the template and target image is desirable.

Given a $M_I \times N_I$ target image $I(x, y)$ such that $I(x, y) = 0$ for $x < 0$ or $x \geqslant M_I$ or $y < 0$ or $y \geqslant N_I$ and a $M_T \times N_T$ template image $T(x, y)$, such that $T(x, y) = 0$ for $x < 0$ or $x \geqslant M_T$ or $y < 0$ or $y \geqslant N_T$, the normalized correlation (NCC) of the target image and template image is written as

$$C(u, v) = \frac{\sum_x \sum_y (I(x, y) - \overline{I(u, v)})(T(x - u, y - v) - \overline{T})}{\sqrt{\sigma_I^2(u, v)\sigma_T^2}}, \tag{1}$$

where

$$\sigma_I^2(u, v) = \sum_{x=u}^{u+M_T-1} \sum_{y=v}^{v+N_T-1} [I(x, y) - \overline{I(u, v)}]^2, \tag{2}$$

and

$$\sigma_T^2 = \sum_x \sum_y [T(x, y) - \overline{T}]^2, \tag{3}$$

$$\overline{I(u, v)} = \sum_{x=u}^{u+N_T-1} \sum_{y=v}^{v+M_T-1} I(x, y)/(N_T M_T) \tag{4}$$

and

$$\overline{T} = \sum_x \sum_y T(x, y)/(N_T M_T). \tag{5}$$

The NCC is a measure of how well the template shifted by (u, v) resembles the target image. The position corresponding to a good match is obtained by noting the shift values u and v which yield a peak value in $C(u, v)$. It can be shown that the NCC provides a measure of similarity between the template signal and the target signal that is independent of changes in signal scaling or signal offset. The NCC has a value $C \in [-1.0, 1.0]$. If the target image

is $I(x, y) = aT(x - x_0, y - y_0) + b$, then the NCC operation between $I(x, y)$ and $T(x, y)$ would yield a maximum for C at (x_0, y_0) such that $C(x_0, y_0) = 1.0$ independent of a and b.

The NCC as expressed in equation (1) is computationally intensive. The numerator is computed via a fast Fourier transform but the term $\sigma_I^2(u, v)$ which measures the variation of the image under the template must be recomputed each time the template moves to a new position. The use of sum tables to compute this term greatly improves performance (Lewis 1995) such that most of the time is spent computing fast Fourier transforms. Details can be found in the appendix.

We have also performed image registration using a simple least-squares measure (LSM) of similarity between the template and the target image defined by

$$LSM(u, v) = \sum_x \sum_y (I(x, y) - T(x - u, y - v))^2 \tag{6}$$

to compare against the NCC-based approach.

2.2. Lesion stabilization

The NCC is used to compare the Hth image in a cine run denoted by $I_H(x, y)$ with a template extracted from the Gth image $I_G(x, y)$, and denoted by $T_G(x, y)$. The Gth template is a rectangular sub-region of $I_G(x, y)$ containing the feature we wish to stabilize. The NCC used for this comparison is denoted by $C_{T_G, I_H}(u, v)$ and yields a shift vector $\vec{S}_{GH} = (u_{GH}, v_{GH})$ used to display the image in the frame of reference of the matched feature. This process is repeated for all frames to be stabilized.

In this paper, different techniques are presented to improve stabilization using the NCC. Each technique described in the next sections is subsequently evaluated with clinical angiographic image sequences using an objective pass/fail criterion.

2.2.1. Shift restrictions. The correlation space $C(u, v)$ can exhibit several local maxima, each providing a different shift vector. As the template is reduced in size or the feature of interest exhibits a small signal relative to background structures, the NCC may exhibit several local maxima having comparable values. As an extreme example, it can be shown that a two-pixel template is a perfect match at every location in an image within a scaling factor and offset. Therefore, a moving feature of interest such as an arterial narrowing may be stabilized successfully and appear stationary over a short sequence of frames but move abruptly if some other local maximum in the correlation space becomes the global maximum in some later frame.

The question of whether it is possible to reduce the number of template matching errors by using physical knowledge of the motion was investigated. Heart motion is cyclical and continuous with a maximum peak motion velocity. Therefore, a constraint was introduced to limit the maximum interframe template shift distance (Sussman and Wright 2003).

For the frame I_H to be stabilized, C_{T_G, I_H} is computed over a sub-region of I_H. This subregion is a super-set of the sub-region defining T_G from I_G but larger by $2\Delta P$ extra pixels along the x and y directions. Thus, interframe shifts $C_{T_G, I_{G+1}}$ greater than $\pm\Delta P$ are not permitted.

2.2.2. Template updating. Stabilization of the second frame in a sequence of images involves the computation of C_{T_1, I_2} using a template extracted from the first frame. Starting with the third and subsequent frames, there exist several candidate frames from which a template can be extracted. The first template could always be used resulting in the sequence of NCC

calculations $C_{T_1,I_2}, C_{T_1,I_3}, \ldots, C_{T_1,I_H}$. This approach is only satisfactory if the appearance of the feature does not change over the sequence. However, coronary motion is the result of myocardial contraction, a process causing deformation of the walls of the heart. As a result, a vessel segment experiences three-dimensional rotation and translation. This results in a change in the appearance of the vessel projection in addition to the effect of *local* deformation, which also causes some bending and twisting of the segment of interest. As a result, it can be advantageous to allow the template to evolve such that for a given image, it is obtained from the previous image resulting in the sequence $C_{T_1,I_2}, C_{T_2,I_3}, \ldots, C_{T_{H-1},I_H}$ thus allowing for changes in the appearance of the feature.

2.2.3. Image equalization. In coronary angiography, the features of interest are primarily contrast filled blood vessels and not the large scale intensity variations due to *other structures* such as the lungs and diaphragm. The stabilization method should also perform better if the signal from the arteries dominates the image. In theory this could be accomplished through the use of DSA. Unfortunately, DSA is very difficult to perform in coronary angiography because of organ motion that occurs between the acquisition of the pre-contrast subtraction mask and the opacified frames causing image artefacts in the subtracted images.

Our approach to reducing signals from other structures involves computing a mask image from each frame through spatial blurring (Robert *et al* 2002). We have shown that spatial convolution of an angiographic frame with a circularly symmetric Gaussian kernel having a standard deviation (characteristic width α) greater than the diameter of the arteries results in an image where these arteries are nearly invisible. Thus, this blurred image can be used as a subtraction mask that is applied to the original image. There are no mis-registration artefacts because the mask is derived from the original image.

2.2.4. Template rotation. A shift (u', v') of the features in the target image I only causes a shift in C. In other words $C(u-u', v-v') = NCC[T(u, v), I(u-u', v-v')]$, and the NCC operator is said to be shift invariant. However, the NCC operator is not rotationally invariant. If the feature of interest in T is present but rotated in I, the NCC value may be reduced. Let us define

$$C_\theta(u, v) = NCC[T_\theta(u, v), I(u, v)], \tag{7}$$

to indicate that C is computed for a number of templates rotated over a range of angles $(\theta_{-n}, \theta_{-n+1}, \ldots, 0, \ldots, \theta_n)$.

The user identifies the template as before but now several rotated templates $T_\theta(u, v)$ are derived from the original template specified by the user. The axis of rotation of the rotated templates corresponds to the centre of the original template. The rotated templates have the same width and height as the original. Under this scheme, the best template matches are obtained by finding the values of θ, u and v corresponding to maxima in $C_\theta(u, v)$.

3. Method

An angiographic viewing programme was developed to demonstrate and evaluate our stabilization approach. The viewer allows the user to scroll through the sequence of images and identify the range of frames to be stabilized. The feature to be stabilized is a rectangular region specified directly on the image with the aid of a computer mouse. The viewer utilizes the Papyrus toolkit (University of Geneva) to read DICOM image files acquired from clinical angiographic suites. All images used in this paper were 512 × 512 pixels by

8 bits. Computations to stabilize the lesion are performed at a rate of 35 frames s^{-1} on a 3.0 GHz Pentium IV computer. Therefore, the entire operation can be done in a few seconds.

3.1. Optimizing stabilization with clinical images

3.1.1. Developing a performance criterion. Forty-five anonymized angiographic runs were selected as a training set for our method from a larger set of clinical angiograms because they exhibited one or more lesions with a $\geqslant$40% diameter stenosis in one or more of the major coronary arteries in which catheter-based interventions would be considered. These included the right coronary, left anterior descending and circumflex arteries. Subsets of frames encompassing one optimally opacified cardiac cycle were identified in each cine run. The initial and final frame indices were recorded.

In order to optimize the stabilization method, we developed an objective criterion for success of the method using a two-step process. First, stabilization was performed on the subsets of frames from 45 cine runs and for each instance, two observers were asked to visually assess whether the method was successful. Sequences where the arterial lesion exhibited one or more abrupt changes in position in adjacent frames were considered as failures. For all the cases examined there was 100% agreement between both observers with regards to which sequences were successfully stabilized.

Secondly, an objective criterion to capture the observer's visual pass/fail criterion was needed. This criterion must be a measurement that does not involve visual estimation or other user interaction so that it can be automated. The 45 sequences were all viewed one frame at a time and for each frame two observers independently recorded the position of the lesion with the aid of a mouse. Both observers were required to place the mouse on the feature located near the narrowest part of the lesion that could be identified with confidence over all of the frames in each sequence. Vessel segments immediately proximal and distal to a lesion are easily visualized and can be used if needed to infer lesion motion due to their immediate proximity to the lesion if the lesion is hard to see. Each user viewed their own cine runs stabilized with their measurements and corrected the shifts where needed repeating the process until satisfied with their results.

Figure 1 shows the distribution of the resulting lesion interframe shift distances obtained from 653 frames. The smaller values correspond to end-systole and end-diastole where motion between successive frames is least. The mean interframe shift distance is 9.4 pixels. Figure 2 shows the distribution of the differences in interframe shift distances between the two observers. For each cine run, both observers together witnessed a version stabilized with their own and the other person's shift measurements and chose the best one.

Automated lesion stabilization was performed on all of the cine runs. Figure 3 shows the distribution of the differences or errors in interframe shift distances for manual and automated stabilization. It was found that cine runs where this error for any adjacent frame pairs was greater than 8 pixels corresponded to the failures first identified visually. Thus, a success/failure criterion that can be automated and is consistent with the observer's definition was developed. This criterion was used to optimize the stabilization method allowing the measurement of the registration accuracy in the absence or presence of the various features described in section 2.2 with the training set of 45 cine runs.

3.1.2. Measuring registration performance with random clinical images. Finally, 21 angiographic studies each containing several cine runs of patients undergoing percutaneous coronary interventions (PCI) were selected at random without viewing any cine runs. For each study, one to three cine runs showing one or more lesions selected for PCI were included for

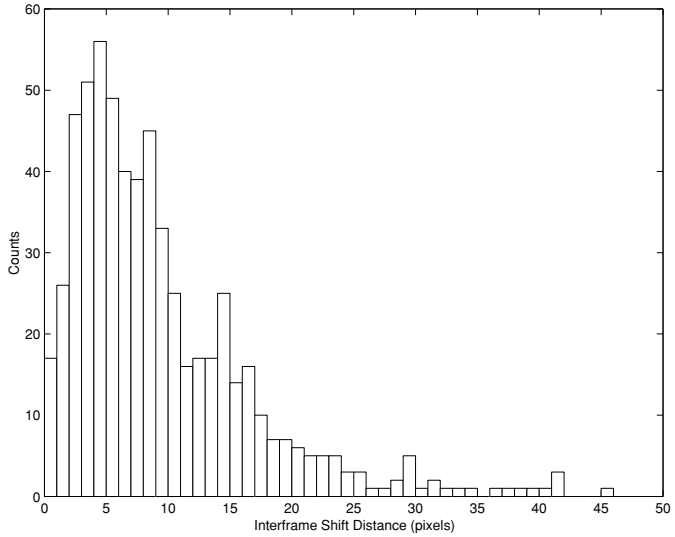

Figure 1. Interframe shift distances of lesions in pixels from 653 frames obtained from 45 cine runs.

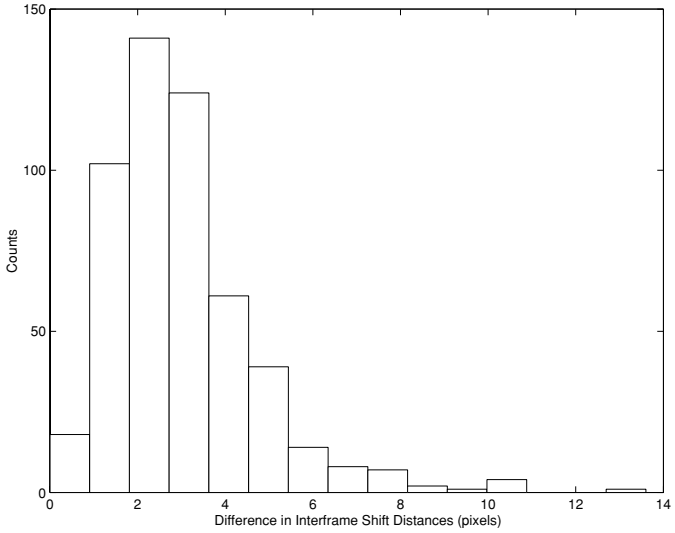

Figure 2. Differences in interframe shift distances of lesions in pixels for two observers from 653 frames selected from 45 cine runs.

a total of 41 cine runs totalling 617 frames. No study was excluded upon examination of the images so all studies contributed at least one cine run. In a few studies, up to 2 cine runs of the same lesion shown from different angles were selected. A few studies exhibiting several lesions undergoing treatment contributed up to 3 cine runs showing distinct lesions. The stabilization method optimized with the training set was applied to these cine runs. Manual stabilization was performed by identifying the lesions in each frame in the same manner as described for the training set in section 3.1.1. Thus, errors in interframe shifts between automated and manual stabilization could be measured and the success rate obtained using the pass/fail criterion described in the previous section.

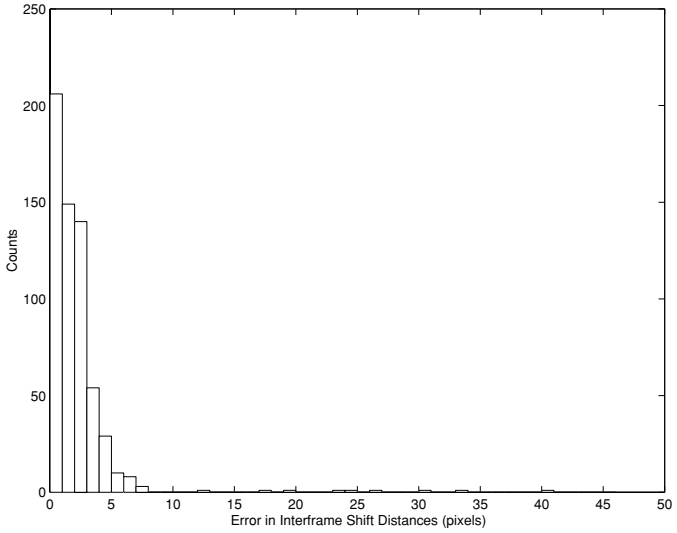

Figure 3. Error in interframe shift distances of lesions in pixels between automated and manual stabilization obtained from 653 frames selected from 45 cine runs.

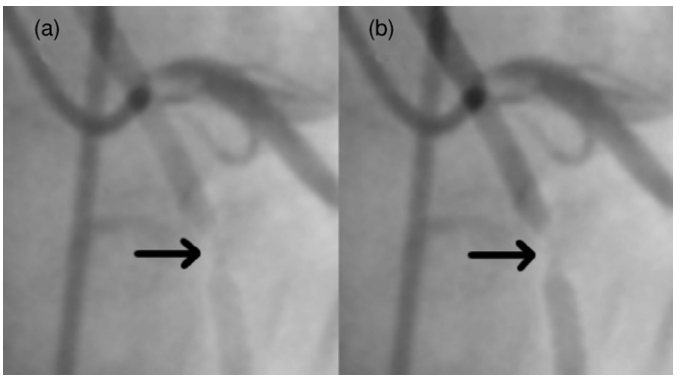

Figure 4. Two clinical images showing a catheter inserted into the left main coronary artery. A simulated stenotic vessel having a diameter similar to real vessels has been added to both images. The arrows point to the stenoses of the simulated vessels. (a) The simulated vessel with the least amount of contrast used in our experiment. (b) The simulated vessel with the highest contrast.

3.2. Performance with simulated vessels

In order to better understand the modes of failure of the stabilization process, a simulated artery was added to clinical angiograms. This simulation allows the quantification of small registration errors because the true interframe motion is known. The artery cross-sections were obtained by projecting a vessel segment having a circular cross-section. The simulated artery was blurred to approximate degradation due to the imaging system creating an artificial artery image $A(x, y)$ which was combined with the clinical image. The diameter was chosen to be similar to that of the major coronary arteries seen in clinical angiograms. Figure 4(a) shows the simulated artery with the smallest contrast used in our experiment. Figure 4(b) shows the largest contrast used in our simulations. This contrast was chosen to yield a simulated vessel similar to real vessels. Arteries having low, intermediate (not shown in figure 4) and

Table 1. Table of experimental results. This table highlights the constraints that were imposed on the stabilization algorithm for each experiment and the total number of cases that were successful.

Experiment	Algorithm	Shift restriction	Template updating	Equalization	Rotation (°)	Success (per cine run)	Success (per frame)
1	NCC	No	No	No	0	10/45	508/653
2	NCC	Yes	No	No	0	23/45	575/653
3	NCC	Yes	Yes	No	0	33/45	624/653
4	NCC	No	No	$\hat{\alpha}$ 10/512	0	15/45	548/653
5	NCC	Yes	Yes	$\hat{\alpha} = 10/512$	0	40/45	644/653
6	LSM	Yes	Yes	No	0	24/45	592/653
7	LSM	Yes	Yes	$\hat{\alpha} = 10/512$	0	37/45	633/653
8	NCC	Yes	Yes	$\hat{\alpha} = 10/512$	±10	36/45	643/653
9	NCC	Yes	Yes	$\hat{\alpha} = 10/512$	±5	40/45	647/653

high contrast were added to 20 different angiographic sequences. The simulated arteries were moved by 20 pixels in both the x and y directions per frame. A total of 450 shift measurements were obtained. For each measurement, the error distance between the true shift and the shift estimate from stabilization was computed.

A measure of the strength of the simulated artery signal relative to the background signal over the template area for the best shift $(\hat{u}, \hat{v})$ was defined as

$$S = \sigma_A(\hat{u}, \hat{v})/\sigma_I(\hat{u}, \hat{v}), \tag{8}$$

where

$$\sigma_A^2(u, v) = \sum_{x=u}^{u+M_T-1} \sum_{y=v}^{v+N_T-1} [A(x, y) - \overline{A(u, v)}]^2, \tag{9}$$

and $\sigma_I^2(u, v)$ was given in equation 2. This signal strength S was computed for all 450 stabilized frames. Thus, the dependence of shift errors on the signal strength can be analysed and the modes of failures of our approach with decreasing signal strength can be investigated.

4. Results and discussion

4.1. Optimizing stabilization with clinical images

The features of our method such as equalization, restriction of the region where the template is matched, the use of the NCC as a measure of similarity and whether the template is updated or rotated were introduced because it was hypothesized that they would improve the registration accuracy. The effect of each of these features on performance was verified by stabilizing the 45 image sequences described previously with and without a given feature and applying the pass/fail criterion described in section 3.1.1.

Table 1 summarizes the results obtained for all of our experiments where the performance of our approach was tested using the 45 clinical sequences of images. Experiment 1 had the lowest success rate. There were only 508/653 frames where the differences in interframe shift distances (between manual and automated stabilizations) were equal to or less than 8 pixels. There were only 10/45 cine run where all frames meet this criterion for success.

The template matching area included the whole image, and the same template used to match all subsequent frames was extracted from the first frame. No equalization or template rotation was performed. Experiment 2 resulted in improved performance through

the introduction of a restriction in the template matching region. Shifts from one frame to the next were restricted to be within ± 32 pixels in the x and y directions.

The improved success rate in experiment 3 confirms the value of template updating. A lesion may have a different appearance at one stage of the cardiac cycle than it does at another stage, so the use of the initial template to identify the lesion in the intermediate frames may result in an incorrect match. Evolving the template after each frame reduced the risk of incorrectly tracking the lesion since frame-to-frame variations between successive frames are small.

Experiment 4 demonstrates the advantage of performing equalization alone prior to stabilization. The cine runs were equalized with mask images obtained by blurring with Gaussian kernels having characteristic width values of 5, 9, 10 ,11, 15, 20 and 30 pixels. The best results were obtained for $\hat{\alpha} = 10(/512)$. In other words, the Gaussian blurring kernel used to produce the subtraction mask image had a standard deviation of 10 pixels for a 512×512 image. The number of successes increased from $10/45$ image sequences (experiment 1) to $15/45$. Feature tracking can fail when lesions move across a region exhibiting strong signals due to background structures. The background variations in the template may contain a lot of signal power, and cause the template matching algorithm to track the background intensity variations rather than the vessel. Equalization reduces the signal due to background structures and thus improves vessel tracking.

Experiment 5 was done with shift restrictions, template updating and equalization combined. The method failed in only $5/45$ cases and only $9/653$ frames were mis-registered.

Experiments 6 and 7 allowed us to evaluate the lesion stabilization method with a simpler least-squares measure (LSM) of similarity between the template and image. The LSM was used to perform template matching instead of the NCC with and without equalization. Experiments 3 and 6 differed only in the use of the LSM as opposed to the NCC to perform stabilization. Use of the NCC resulted in an increase in the number of successes from 24 to 33. Similarly, experiments 5 and 7 differed only in the use of the LSM as opposed to the NCC to perform stabilization with equalization. The LSM approach was successful in 37 versus 40 cases for the NCC.

Experiment 8 was conducted using the best conditions found in experiment 5 with the addition of template rotation described in section 2.2.4. The templates were rotated in 1° increments from -10 to $+10^\circ$ and used for template matching. This increases the computational burden twenty fold. Figure 5 shows the distribution of interframe rotation angles for the 653 frames registered in the 45 cine runs obtained from the data in Experiment 8. The distribution has a standard deviation of 2.8° and shows that interframe rotations are generally small.

Rotation over this range of angles yields a lower success rate as it increases the likelihood of the rotated template matching against some other vessel segment having a different orientation. If the range of angles is restricted to $\pm 5^\circ$ as shown in experiment 9, the success rate becomes comparable with experiment 5. Rotation provides no significant advantage if the template is always picked from the most recent previous frame because interframe rotation angles are small and cause a linear increase in computation time with the number of rotated templates used.

4.2. Measuring registration performance with random clinical images

Stabilization using the optimized conditions described in experiment 5 was applied to the 41 new cine runs obtained from 21 interventional angiographic studies. Stabilization performance was measured using the 8 pixel limit for errors in interframe shift distance

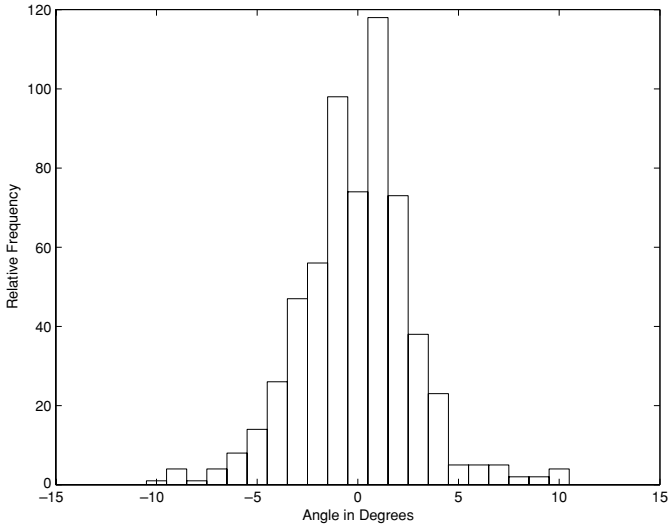

Figure 5. The distribution of interframe rotation angles for 653 frames from 45 cine runs. The interframe rotation angles are generally small as there is little rotation of the lesions of interest between adjacent frames. This plot was obtained using the condition for experiment 8 described in table 1.

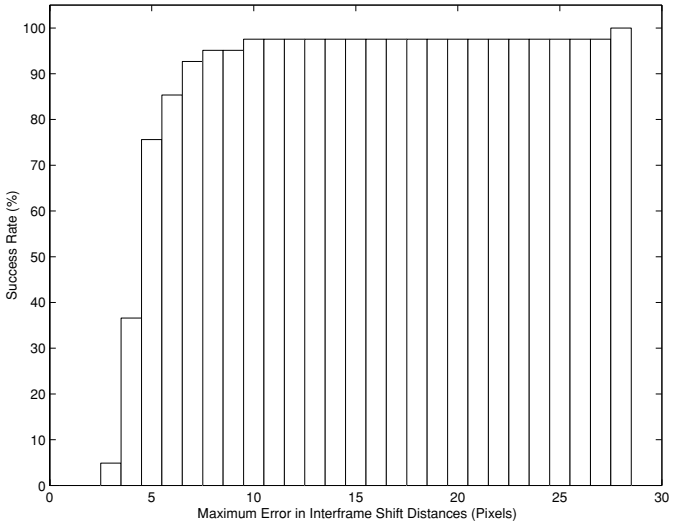

Figure 6. A plot of the success rate for the stabilization of 41 cine runs versus the maximum error in interframe shift distances allowed between manual and automated stabilizations. The stabilization was done using the conditions described in experiment 5 shown in table 1.

as described in section 3.1.1. A total of 39/41 cine runs passed this test. Figure 6 shows the dependence of this success rate as a function of the maximum allowed error in interframe pixel shift.

The template must encompass a region having vertical and horizontal dimensions greater than the vessel diameter in order to track the stenosis. Thus the template provides a measure of the aggregate motion of the lesion weighted towards the part of the vessel having the strongest

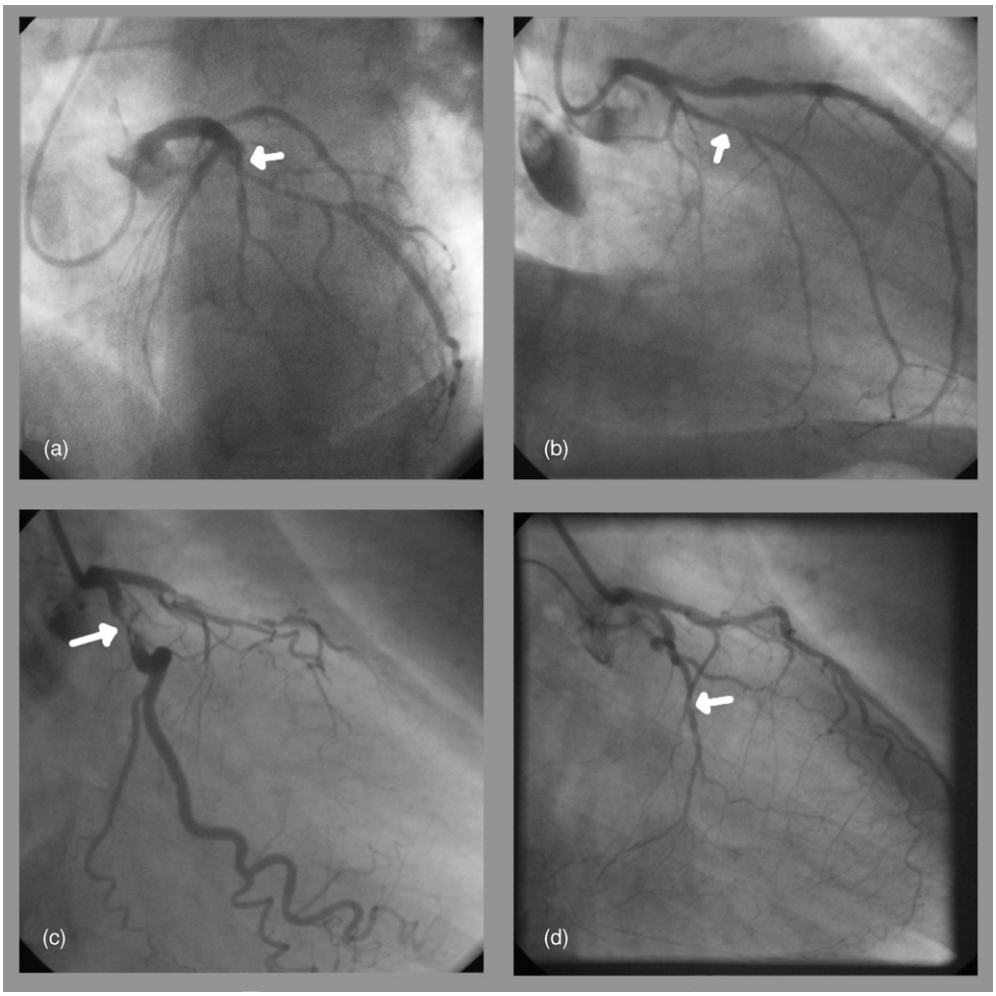

Figure 7. Examples of stabilized cine runs that failed or nearly failed according to our success criterion based on a comparison with manual stabilization. Panels (a) and (b) show the only two cine runs where stabilization failed out of 41 cine runs. Panels (c) and (d) show the two cine runs that were closest to failure.

signal and not necessarily the exact same point specified with the mouse when doing manual stabilization. The deformation of the lesion as a function of the cardiac cycle also makes it difficult to manually define exactly the same point for all frames. For these reasons, there are no cine runs with errors in interframe shift distances less than 3 pixels.

Figure 7 shows in panels (a) and (b) frames from the only two cine runs that failed according to our automated success criterion. Stabilization in (a) failed because one frame barely exceeded the maximum allowed difference in interframe shifts between manual and automated stabilizations due to interference caused by an overlapping vessel. Stabilization in (b) failed because of the presence of some background structures making the lesion appear less significant in one frame. The two cine runs in panels (c) and (d) were closest to failing but passed according to said criterion. Overlapping vessels made the lesions difficult to see for some frames.

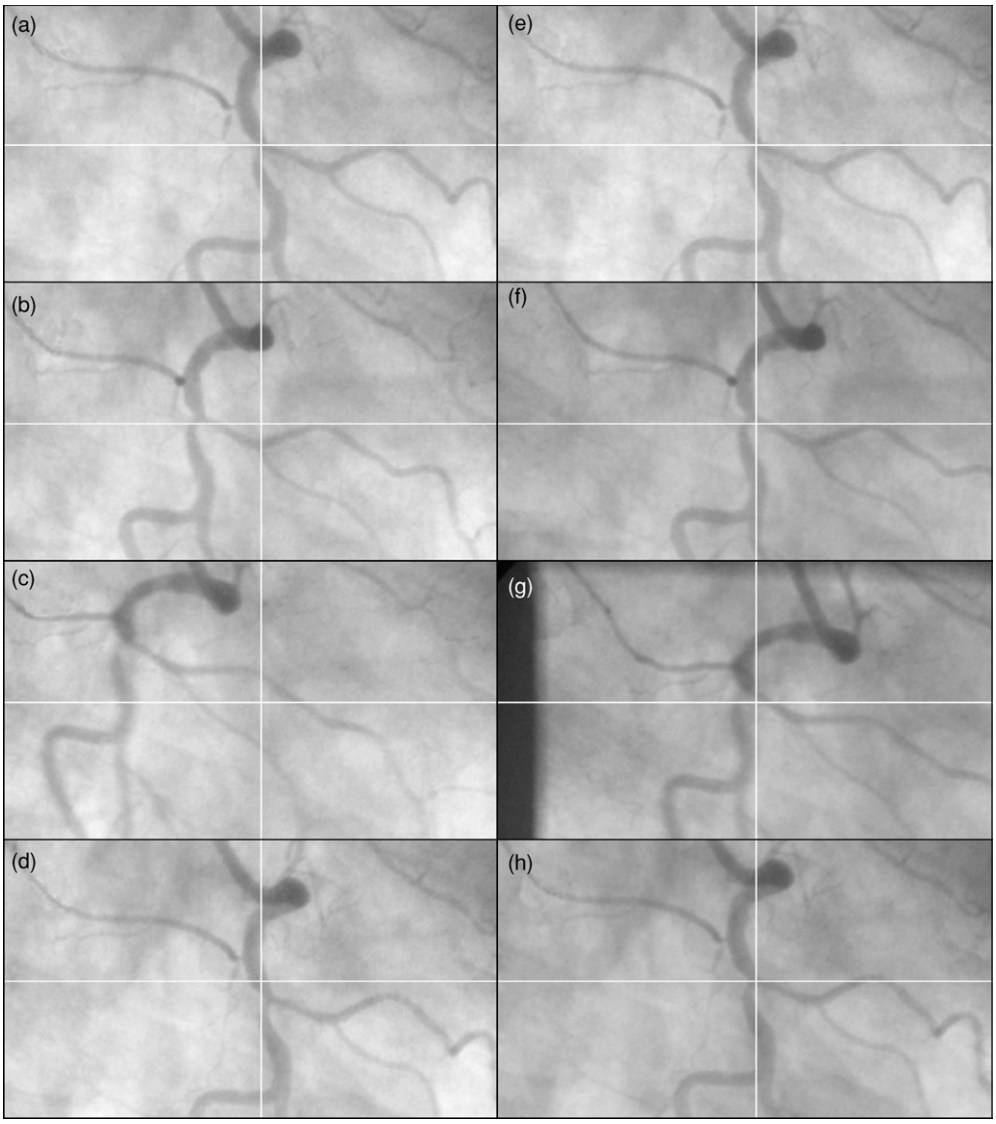

Figure 8. Images (a)–(d) show the first, fourth, sixth and twelfth frames in a sequence of angiograms. Frame twelve is in the same phase of the cardiac cycle as the first frame. The position of the artery varies due to cardiac motion. Cross-hairs are added in the same position in each image to help assess the motion of the stenosis. Images (e)–(h) show the frames following stabilization. The stenosis only moves by a small amount in intermediate frames shown in (f) and (g) and this motion is easily accommodated by the observer. The lesion is in the same position in the last frame as it was in the first frame so the sequence images could easily be viewed in a loop without any motion discontinuities.

4.3. Viewing stabilized frames

Typically, stabilization is done with the most opacified frames because they exhibit the highest signal-to-noise ratio. These frames encompass one or two cardiac cycles and only last 1 to 2 s. The user will likely want to see these frames displayed continuously without intervention over

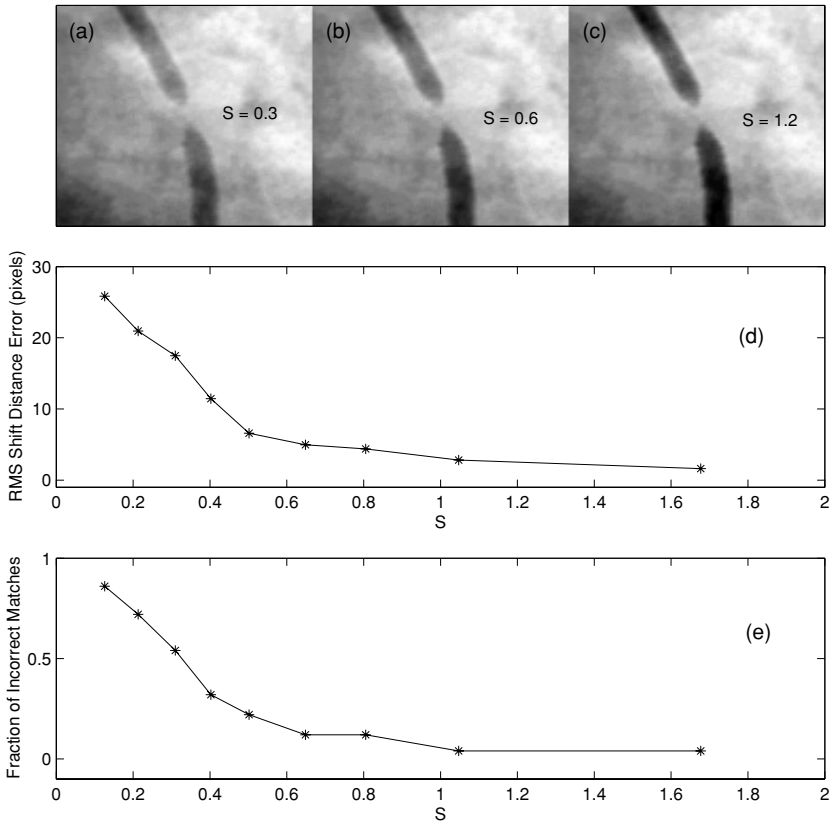

Figure 9. Stabilization of simulated vessels added to clinical images. (a)–(c) Simulated vessels with differing artery signal strength relative to the background signal denoted by S. (d) The root mean squared shift distance error is plotted as a function of S. The error measurement is the distance between the estimated location of the vessel and its true location (see text). (e) A plot of the fraction of incorrect shift estimates (shift distance error >2 pixels). Each point in both plots is computed from 50 shift error distance measurements having similar S values.

and over in a loop. In that case, the arterial segment of interest may suddenly change in appearance at the transition from the last frame back to the first frame making the sequence of images difficult to watch. This problem is most easily addressed by showing the frames in forward and then reverse order continuously.

If the user prefers to view the frames in a loop in their true temporal order, motion discontinuities between the last and first frames can be minimized (Langer and Robert 2003). Briefly, the choice of last frame is refined to be in the same phase of the cardiac cycle as the first frame. This is achieved by finding the last frame yielding the best $C(u, v)$ value with the first frame. If the template is allowed to evolve as described in section 2.2.2 some residual motion discontinuities may still occur because the manner in which the template evolves cannot be predicted *a priori*. In that case, the shift value (u, v) between the first and last frames is distributed linearly over all frames in the sequence to eliminate this residual discontinuity (Langer and Robert 2003). The addition of these corrections is transparent to the user because it requires no additional input and adds less than 1 s to the total processing time.

Figure 8 shows from top to bottom the first, fourth, sixth and twelfth frames of a cine run without stabilization on the left and with stabilization on the right. Cross-hairs are overlayed

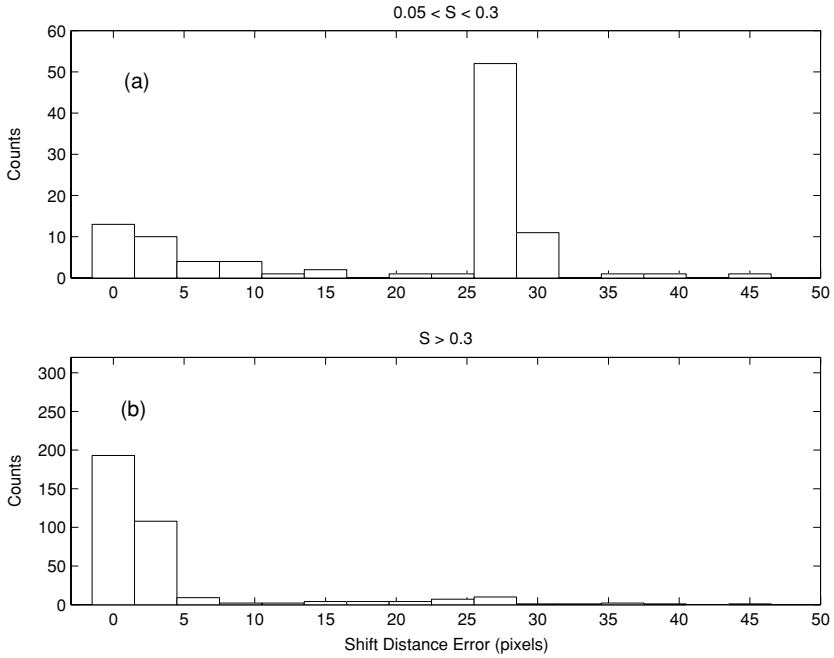

Figure 10. Histograms showing the errors in determining the position of a simulated arterial segment added to clinical angiograms. The error measurement is the distance between the estimated location of the vessel and its true location. The histograms summarize 450 attempts to measure the shift of the simulated vessel. (a) In the low S (artery signal strength relative to the background) case, stabilization often failed leading to large errors when the template matched the background instead of the moving vessel as illustrated by the peak for error shifts around 30 pixels in (a). (b) In the higher S case most shift measurements were correct.

on the images of a stenosis just above an arterial bifurcation in the first frames (a) and (e) to make the motion of this feature easier to assess. In subsequent frames, the cross-hairs which are fixed highlight the motion of the stenosis in the unstabilized frames and the extent of the problem.

The twelfth frame in the sequence shown in (d) was chosen as the last frame because it was in the same phase of the cardiac cycle as the first frame shown in (a). If the patient and x-ray detector did not move, the arteries in (a) and (d) should be roughly in the same position with small differences due to respiratory motion whose period is independent of the length of the cardiac cycle.

Figures 8(e)–(h) show the same frames as in (a)–(d) following stabilization. Closer examination of the bifurcation for intermediate frames (f) and (g) shows a small drift in the artery relative to the cross-hairs which can easily be accommodated by the observer. This is due to changes in the configuration of the arteries during the cardiac cycle.

However, the cross-hairs in the last frame are at the same position as in the first frame due to the successful application of our last to first frame stabilization method.

It is worth noting that another source of discontinuity between the last frame and the first frame due to changes in contrast and brightness may occur. However, this grey level discontinuity is generally small and it can be corrected by manipulating the brightness of the images in the sequence.

4.4. Performance with simulated vessels

A total of 450 pairs of shift error distance measurements and S values were obtained. Figures 9(a)–(c) show simulated vessels with differing S values as a reference. Figure 9(d) shows the root mean squared (RMS) shift distance error plotted as a function of S. The extreme left point on the graph was obtained by calculating the RMS distance error for the 50 smallest S values. Each subsequent point was calculated from the next 50 smallest S values. Figure 9(e) shows the fraction of shift estimates that are in error by more than two pixels. Each point was also obtained by combining 50 experimental results having similar S values as in (d). The fraction of tracking failures rose rapidly for S values less than 0.6.

Histograms of the distribution of shift errors for the low and high S cases were also computed and are shown in figure 10. At high S values, the shift errors were nearly always zero indicating that the vessel was tracked correctly. At low S values, non-zero shift errors were more common but were not randomly distributed. This is because signals due to background structures start to dominate the image and a match is made on these features instead of the artery. The simulated vessel moved a distance of 20 pixels in the x and y directions between subsequent frames so the tracking error should be approximately 28 pixels if the NCC matches the background structures which are stationary or moving slowly (image intensifier panning or patient motion). The low S value histogram shows a peak in the distribution of shift errors around 30 pixels. The distribution of tracking errors over all S values is bimodal as the method either matched the moving feature of interest or background structures.

5. Conclusion

We have shown that lesion stabilization can be automated such that it can be performed on a feature of interest in a few seconds. The introduction of restrictions on the motion of the template, equalization and template updating all improved the success of our approach. The success rate of our optimized method when applied to 41 clinical angiographic sequences was 95% (39/41). Experiments where moving simulated vessels were added to clinical images showed that the registration method either tracks the feature of interest accurately or for low vessel signal strength it may lock onto the background structure in the image.

Stabilizing the lesion allows the observer to maximize the ability to temporally average successive frames in a sequence improving the perceived SNR of the lesion, and allows one to monitor the lesion over the entire cardiac cycle without any distractions. Failures in the registration procedure are manifested as a sudden jump in the position of the lesion and are obvious. As a result, there is no risk of altering the diagnostic information in an imperceptible way.

Appendix

Here we describe the fast NCC approach (Lewis 1995). The description is more detailed than its original form with respect to the use of sum tables and is provided for the convenience of the reader.

The NCC as expressed in equation (1) is lengthy to calculate. If we define $T^a(x, y) = T(x, y) - \overline{T}$, the numerator of equation (1) can be rewritten as

$$\sum_x \sum_y I(x, y) T^a(x - u, y - v). \tag{A.1}$$

This constitutes a two-dimensional correlation operation which can be calculated rapidly in Fourier space.

The second term inside the square root, $\sum_x \sum_y [T(x, y) - \overline{T}]^2$, measures the variation in the template image. The term is independent of u and v and can be pre-computed for a given template. The term $\sum_{x=u}^{u+M_T-1} \sum_{y=v}^{v+N_T-1} [I(x, y) - \overline{I(u, v)}]^2$ measures the variation of the image under the template and requires numerous calculations because $\overline{I(u, v)}$ must be determined for all values of the template position variables u and v. This term can be computed far more efficiently by using pre-computed sum tables (Lewis 1995). For an $M_I \times M_I$ pixel target image and $M_T \times M_T$ pixel template image, it can be shown that the number of operations required to compute the image variation factor is reduced from $[M_I(M_I - M_T + 1)]^2$ to $[4(M_I - M_T + 1)]^2$ operations through the use of these tables. Equation (A.4) shows the NCC in its final form where the numerator is computed in Fourier space and the first term of the denominator is determined with the aid of sum tables. In this form, most of the computation time is spent calculating fast Fourier transforms.

The sum table S_1 is defined as

$$S_1(u, v) = \begin{cases} \sum_{x=0}^{u} \sum_{y=0}^{v} I(x, y) & \text{if} \quad u \geqslant 0 \text{ or } v \geqslant 0 \\ 0 & \text{otherwise} \end{cases}, \tag{A.2}$$

and the sum table S_2 is defined as

$$S_2(u, v) = \begin{cases} \sum_{x=0}^{u} \sum_{y=0}^{v} I(x, y)^2 & \text{if} \quad u \geqslant 0 \text{ or } v \geqslant 0 \\ 0 & \text{otherwise} \end{cases}. \tag{A.3}$$

Using the sum table and the Fourier transform, the NCC can be written as

$$C(u, v) = \frac{\mathcal{F}^{-1}\{\mathcal{F}\{T^a(x, y)\}^* \mathcal{F}\{I(x, y)\}\}}{[\Delta S_2(u, v) - \Delta S_1(u, v)^2 / [M_t N_T]] \sum_x \sum_y [T(x, y) - \overline{T}]^2} \tag{A.4}$$

where $\mathcal{F}$ is the two-dimensional Fourier transform, $\mathcal{F}^{-1}$ is the inverse two-dimensional Fourier transform, the * symbol represents the complex conjugate operator,

$$\Delta S_1 = S_1(u + N_T - 1, v + M_T - 1) - S_1(u - 1, v + M_T - 1) \tag{A.5}$$

$$- S_1(u + N_T - 1, v - 1) + S_1(u - 1, v - 1), \tag{A.6}$$

and

$$\Delta S_2 = S_2(u + N_T - 1, v + M_T - 1) - S_2(u - 1, v + M_T - 1) \tag{A.7}$$

$$- S_2(u + N_T - 1, v - 1) + S_2(u - 1, v - 1). \tag{A.8}$$

References

Achenbach S, Ropers D, Holle J, Muschiol G, Daniel W G and Moshage W 2000 *Radiology* **216** 457–63

Berthelsen B and Cederblad A 1991 *Acta Radiol.* **32** 492–7

Cox J S and de Jager G 1994 *Proc. SPIE* **2167** 188–99

Eckstein M P, Whiting J S and Thomas J P 1996 *Proc. SPIE* **2712** 9–25

Eigler N L, Eckstein M P, Mahrer K N and Whiting J S 1994 *Circ.* **89** 2700–9

Fitzpatrick J M, Grefenstette J J, Pickens D R, Mazer M and Perry J M 1988 *Image Processing in Medical Imaging* ed C N de Graaf and M A Viergever (New York: Plenum) pp 415–35

Heart and Stroke Statistical Update 2001 Technical report American Heart Association Dallas, Texas

Hofman M B M, Wickline S A and Lorenz C H 1998 *J. Magn. Reson. Imaging* **8** 568–76

Langer A and Robert N 2003 US Patent 6496716 Method and apparatus for stabilization of angiography images

Levin D C and Gardiner G A 1992 *Heart Disease* ed E Braunwald (Philadelphia: Saunders) chapter 36 264

Lewis J P 1995 *Proc. Vision Interface 95 (Québec)* pp 120–3

Meijering E H, Niessen W J and Vierger M A 1999 *IEEE Trans. Med. Imaging* **18** 2–21

Robert N, Komljenovic P T and Rowlands J A 2002 *Med. Phys.* **29** 736–47
Rosenthal L *et al* 1998 *Am. J. Cardiol.* **82** 451–8
Sekuler R and Blake R 1990 *Perception* 2nd edn (New York: McGraw-Hill)
Sussman M S and Wright G A 2003 *IEEE Trans. Med. Imaging* **22** 206–16
Wang Y, Vidan E and Bergman G S 1999 *Radiology* **213** 751–8
Wexler L 1995 *A Categorical Course in Physics. Physical and Technical Aspects of Angiography and Interventional Radiology* ed S Balter and T B Shope (Chicago: RSNA Publications)
Whiting J S, Honig D, Carterette D and Eigler N L 1991 *Proc. SPIE* **1453** 165–7
Wilson D L, Xue P and Aufrichtig R 1994 *Med. Phys.* **21** 1875–83
Wright G A, Taylor K W and Rowlands J A 1985 *Med. Phys.* **12** 705–12
Xue P and Wilson D L 1998 *J. Opt. Soc. Am.* A **15** 375–88