

Motion compensated digital subtraction angiography

Magnus Hemmendorff^{ab}, Hans Knutsson^a
Mats T. Andersson^a, Torbjörn Kronander^b

^aLinköping University, Electrical Engineering, SE-581 83 Linköping, Sweden

^bSECTRA, Teknikringen 2, SE-583 30 Linköping, Sweden

ABSTRACT

Digital subtraction angiography, whether based on traditional X-ray or MR, suffers from patient motion artifacts. Until now, the usual remedy is to pixel shift by hand, or in some cases performing a global pixel shift semi-automatically. This is time consuming, and cannot handle rotations or local varying deformations over the image. We have developed a fully automatic algorithm that provides for motion compensation in the presence of large local deformations. Our motion compensation is very accurate for ordinary motions, including large rotations and deformations. It does not matter if the motions are irregular over time. For most images, it takes about a second per image to get adequate accuracy. The method is based on using the phase from filter banks of quadrature filters tuned in different directions and frequencies. Unlike traditional methods for optical flow and correlation, our method is more accurate and less susceptible to disturbing changes in the image, e.g. a moving contrast bolus. The implications for common practice are that radiologists' time can be significantly reduced in ordinary peripheral angiographies and that the number of retakes due to large or local motion artifacts will be much reduced.

Keywords: Motion Compensation, Motion Estimation, Digital Subtraction Angiography, Aperture Problem, Pixelshift, Automatic, Quadrature Filter, Regularization

1. INTRODUCTION

Cardiovascular disease is the leading killer throughout the industrial world. According to U.S. Department of Health and Human Services¹, more than 950,000 Americans die of cardiovascular disease each year, accounting for more than 40% of all deaths. About 57 million Americans, nearly one fourth of the U.S. population, live with some form of cardiovascular disease.

The standard method of diagnostics is angiography, i.e. medical imaging on vasculature. Angiographies can be acquired using conventional X-ray, CT, MR or other image capturing device. MR angiography can be performed using time of flight (TOF) or phase contrast (PC) or with intravenous injections of contrast agents. For X-ray and CT, there is no alternative to using contrast agents. In digital subtraction angiography, a contrast bolus is injected meantime a sequence of images is acquired. Despite contrast agents, it may be difficult to see the vasculature among other structure in the image, such as bones and variations in tissue. Therefore image subtraction is used on peripheral angiographies, i.e. all images in the sequence are subtracted by an image that was acquired before the contrast bolus reached the region of interest. In case there are no patient motions, only the vasculature remains in the difference images.

Unfortunately, patients often feel a burning sensation of the contrast injection and moves a little. This causes artifacts in the subtraction images. Until now, the usual remedy is to pixel shift by hand, or in some cases performing a global pixel shift semi-automatically. This is time consuming, and cannot handle rotations or local varying deformations over the image. Patient motions cause retakes, where the patient is exposed to extra X-dose and injections of contrast agent, that might be harmful. Contrast agents for X-ray sometimes cause idiosyncratic reactions and chronic kidney damages.

In order to do automatic motion compensation, we need a method to estimate motions. We think that a good motion estimation algorithm for this application should satisfy the following criteria:

Further author information: (Send correspondence to M.H.)

M.H.: E-mail: ma-hem@sectra.se

H.K.: E-mail: knutte@isy.liu.se

- account for irregular motions due to the low frame rate and jerky patient motions.
- not be sensitive to variation of brightness variation between frames, that are caused by tiny variations of X-ray dose.
- not be sensitive to the motion of the contrast bolus.
- not depend on corners and line crossings to solve the aperture problem – often there are only smooth edges in a medical image.
- generate a motion estimate for every point in the image.

We also believe that we can achieve better accuracy if we assume there are no motion discontinuities across a single image.

In the next section we will introduce a novel method that locally estimates constraints on the motion in a neighborhood. Section 3 describes how to fit a motion estimate to these local constraints.

2. PHASE-BASED MOTION ESTIMATION

Using quadrature filters phase is a relatively common approach in stereo algorithms^{2,3}. The idea of using phase for motion estimation has previously been investigated by independent researchers⁴⁻⁶, but to our knowledge, nobody has tried this approach, which extends the accurate stereo algorithms to track motions. Our method is basically a gradient-based method applied after nonlinear preprocessing of the images. To improve accuracy, a confidence measure has been added.

2.1. The quadrature filters

In general, a filter is a quadrature filter⁷ if its Fourier transform, $F(\mathbf{u})$, has zero amplitude on one side of a hyperplane through the origin, i.e. there is a direction $\hat{\mathbf{n}}$ such that

$$F(\mathbf{u}) = 0 \quad \forall \hat{\mathbf{n}}^T \mathbf{u} \leq 0 \tag{1}$$

In two dimensions, a filter is a quadrature filter if its Fourier transform is zero on one half plane. Quadrature filters are closely related to Gabor Filters and also analytic signals. Note that quadrature filters must be complex in the spatial domain. We use a set of four quadrature filters tuned in directions, $\varphi = 0, 45, 90, 135$ degrees. These filters

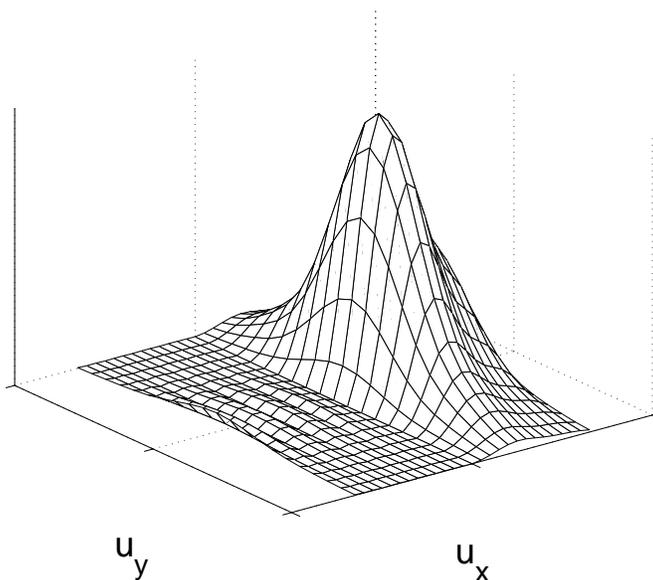


Figure 1. Our quadrature filter in direction $\varphi = 0$. (This filter is not perfect due to the small size of the kernel)

are denoted $f_\varphi(\mathbf{x})$ in spatial domain and $F_\varphi(\mathbf{u})$ in Fourier domain. To be more precise, our filters can be written in the form

$$F_\varphi(\mathbf{u}) = \begin{cases} 0 & \text{if } \hat{\mathbf{n}}_\varphi^T \mathbf{u} \leq 0, \\ F_r(\|\mathbf{u}\|) (\hat{\mathbf{n}}_\varphi^T \hat{\mathbf{u}})^2 & \text{if } \hat{\mathbf{n}}_\varphi^T \mathbf{u} \geq 0, \end{cases} \quad \text{where} \quad \hat{\mathbf{n}}_\varphi = \begin{pmatrix} \cos \varphi \\ \sin \varphi \end{pmatrix} \quad (2)$$

where $F_r(\|\mathbf{u}\|)$ is a real and nonnegative function of the radius in the Fourier domain.

Example: Frequency is the gradient of the phase

We will study what happens if we apply a quadrature filter with direction $\hat{\mathbf{n}}$ on an image where a small neighborhood can be approximated by a single frequency, i.e. $I((x)) = A_{\text{in}} \cos(\mathbf{u}_0^T \mathbf{x} + \theta)$. Without loss of generality, we assume that $\hat{\mathbf{n}}_\varphi^T \mathbf{u}_0 \geq 0$. The filter output will have the form $q((x)) = A_{\text{out}} e^{i \mathbf{u}_0^T \mathbf{x} + \theta}$. The phase of the filter output is $\arg q((x)) = \mathbf{u}_0^T \mathbf{x} + \theta$. This means that the frequency in the image is the slope of the phase. The phase is also linear.

Of course, this approximation is only valid in a neighborhood whose size is as small as the filter kernel. As a justification for using phase in motion estimation, we want to make a comparison with the gradient method for optical flow⁸. We believe that sinusoidal functions provide better approximations than first order Taylor series, which is the approximation in the gradient method.

2.2. Motion Constraint Estimation

We estimate motion of each of the frames individually, compared to a reference image. The intensity of the reference image is denoted $I_A(\mathbf{x})$ and the target image is denoted $I_B(\mathbf{x})$. The quadrature filters are applied in parallel on both of the image frames. The filter outputs for each of the four directions are analyzed separately, in order to avoid structure in different directions to interfere in the motion estimation. The quadrature filters also suppress undesired features like DC value and high frequencies. Unlike conventional gradient methods, our method is not sensitive to low pass variations in image intensity, that are frequent in medical X-ray images, due to random variations in X-ray dose.

For each of the filter outputs, we compute constraints on the local motion. The first step is to convolve image intensities of the two frames, $I_A(\mathbf{x})$ and $I_B(\mathbf{x})$, with quadrature filters

$$q_{A,\varphi}(\mathbf{x}) = (f_\varphi * I_A)(\mathbf{x}) \quad \text{and} \quad q_{B,\varphi}(\mathbf{x}) = (f_\varphi * I_B)(\mathbf{x}) \quad (3)$$

where $\varphi = 0, 45, 90, 135$ and $f_\varphi(\mathbf{x})$ is a quadrature filter.

$$\theta_{A,\varphi}(\mathbf{x}) = \arg q_{A,\varphi}(\mathbf{x}) \quad \text{and} \quad \theta_{B,\varphi}(\mathbf{x}) = \arg q_{B,\varphi}(\mathbf{x}) \quad (4)$$

In all ensuing computations, we must remember that phase is always modulo 2π , but for readability we drop this in our formulas and notations. In most image points, the filter outputs are strongly dominated by one frequency, which makes the phase nearly linear in a local neighborhood. When the phase is linear, it can be represented by its value and gradient. Thus, a gradient method applied on the phase will be very accurate. Of course, the phase is not always linear in a local neighborhood, but that can be detected, and reflected by a confidence measure.

For each point in the image, and for each quadrature filter output, a constraint on the local motion is computed. To simplify notations, we drop the coordinate \mathbf{x} and the direction of the quadrature filter, φ .

$$\mathbf{c} = \begin{pmatrix} c_x \\ c_y \\ c_t \end{pmatrix} = C \begin{pmatrix} \frac{1}{2} \frac{\partial}{\partial x} (\theta_B + \theta_A) \\ \frac{1}{2} \frac{\partial}{\partial y} (\theta_B + \theta_A) \\ \theta_B - \theta_A \end{pmatrix} \quad (5)$$

The motion constraint vector is the spatiotemporal gradient of the phase, weighted by the confidence measure, C , which will be introduced in next section.

2.3. Confidence Measure

Using a confidence measure is necessary to give strong features precedence over weaker features and noise. In addition, it is necessary to avoid phase singularities^{2,9} which occur when two frequencies interfere in the filter outputs. These singularities must be discovered and treated as outliers. All this is done by assigning a confidence value to each constraint. Our confidence measure is inspired by the stereo disparity algorithm by Westelius², which

in turn is inspired by Wilson-Knutsson¹⁰. It is a product of several factors, where the most important feature is the magnitude.

Our confidence measure for magnitude may seem complicated at first glance. Except for suppressing weak features, it is also sensitive to differences between the two frames. This reduces the influence of structure that only exists in one of the images, such as moving contrast bolus and other structure not moving in coherence with the motion we estimate.

$$C_{mag} = \frac{\|q_A\|^2 \|q_B\|^2}{(\|q_A\|^2 + \|q_B\|^2)^{3/2}} \quad (6)$$

Other factors have been added to reflect whether the gradient, is sound for the specific quadrature filter in use. Remind that quadrature filters are only sensitive to frequencies in its direction, not in opposite direction. Thus, the filter outputs should not have any frequencies in opposite direction of the filter direction. In case we still detect negative frequencies, it is because two frequency components interfere. Negative frequencies are illegal and indicate phase singularities^{9,2}.

$$C_{freq>0} = \begin{cases} 1 & \text{if } \hat{\mathbf{n}}^T \nabla \theta > 0, \\ 0 & \text{otherwise,} \end{cases} \quad \text{where } \hat{\mathbf{n}} \text{ is the direction of the quadrature filter in use.} \quad (7)$$

We have also used some more features in our confidence measure, that are omitted here. Those are sensitive to consistency between the two frames and probability of phase wrap around (2π).

2.4. Multiple scales and iterative refinement

To estimate large motions with best possible accuracy, we apply motion estimation iteratively in multiple scales. We begin at the coarsest scale in a low pass pyramid to compute a rough estimate. Then we warp the image, or its filter outputs, and do a new iteration at a finer scale. For best accuracy, we can do multiple iterations at each scale.

When estimating a motion constraints from a warped image, we get a constraint on the motion relative to the warp. Similarly, subsampling alters the estimated motion constraints to yield smaller motion estimates. It is, however, simple to compensate for the warp and subsampling. Assume the image is warped (w_x, w_y) pixels and subsampled λ octaves prior to estimation of a motion constraint, $\tilde{\mathbf{c}} = (\tilde{c}_x, \tilde{c}_y, \tilde{c}_t)$. Then we have in fact estimated that

$$\tilde{c}_x \frac{v_x - w_x}{2^\lambda} + \tilde{c}_y \frac{v_y - w_y}{2^\lambda} + \tilde{c}_t = 0. \quad (8)$$

Thus, the correct motion constraint is

$$\mathbf{c} = \begin{pmatrix} \tilde{c}_x \\ \tilde{c}_y \\ 2^\lambda \tilde{c}_t - \tilde{c}_x w_x - \tilde{c}_y w_y \end{pmatrix}. \quad (9)$$

2.4.1. Warping with integral shifts

Iterative refinement usually involves warping of the image. This requires subpixel interpolation and recomputation of filter outputs and gradients. We have developed a method that applies the warp on filter outputs and gradients. Thanks to compensation of motion constraints, eq. (9), it is not necessary to warp with subpixel accuracy. All rounding errors are cancelled out in the compensated constraints. This scheme works only if rotations and deformations are small.

In case of large local rotations and deformations, it is necessary to map the gradients. For example, a warp that rotates the phase and gradient images must also rotate the direction of the gradient by the same amount.

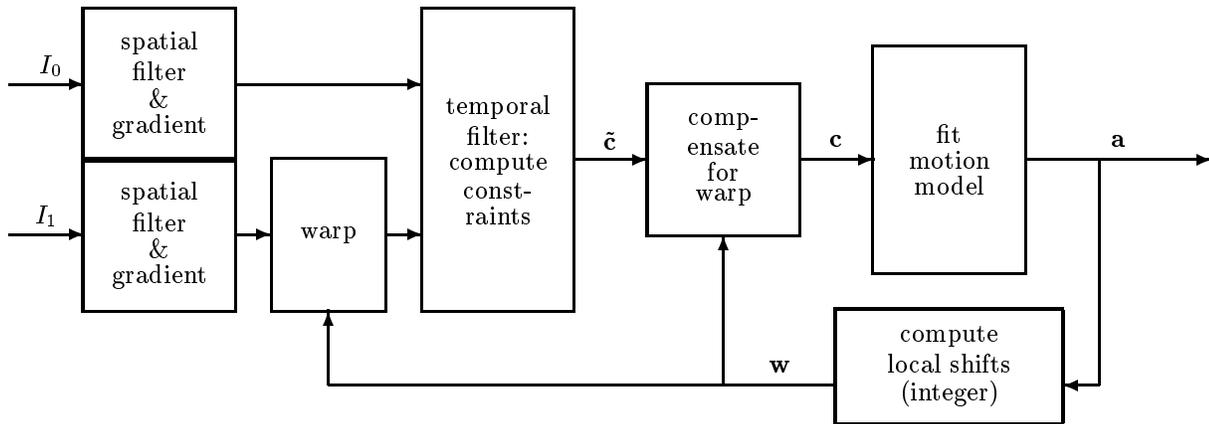


Figure 2. Our scheme of warping. Instead of warping the image, we warp the filter outputs. Since we compensate for the warp directly in the constraints, we do not have to warp with subpixel accuracy.

3. MOTION MODELS

Assume motion constraints have been estimated according to section 2, i.e. we have four motion constraints for each spatial position. The task is to find the motion that best fits the given motion constraints. We perform a least square fit by minimizing an error measure that is defined as the sum of all fitting errors in the entire image,

$$\varepsilon(\mathbf{v}) = \sum_{\varphi=0,45,90,135} \int_{\mathcal{W}} (\mathbf{c}_{\varphi,\mathbf{x}}^T \mathbf{v})^2 d\mathbf{x} = \sum_{\varphi=0,45,90,135} \int_{\mathcal{W}} \mathbf{v}^T \mathbf{c}_{\varphi,\mathbf{x}} \mathbf{c}_{\varphi,\mathbf{x}}^T \mathbf{v} d\mathbf{x}. \quad (10)$$

Our motion estimation is supposed to handle, motions that are not pure translations. This means that \mathbf{v} is not constant over the image. Due to the aperture problem, it is not possible to estimate the motions locally in the image. The aperture problem is usually worse for medical images than real world images, due to lack of corners and line crossings to track. On the other hand, we can assume that all points in the image correspond to the same object and that motions in different parts of the image are related. We do not recommend trying to estimate motions locally in the image and then applying regularization on the motion vectors. Such methods would perform badly since small patches suffer from the aperture problem and local motion estimates would only be accurate in one direction.

3.1. The affine motion model

Three dimensional rigid motions, i.e. translation and rotations, correspond to affine transformations in the image plane. Thus, the affine motion model is adequate for all images where there are no patient deformations. Even in case there are deformations, a local-global variant of affine model is useful. In the affine motion model, the motion is given by

$$\mathbf{v} = \begin{pmatrix} v_x \\ v_y \\ 1 \end{pmatrix} = \begin{pmatrix} a_1 & a_2 \\ a_4 & a_5 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} a_3 \\ a_6 \\ 1 \end{pmatrix} \quad (11)$$

where a_1, a_2, \dots, a_6 are parameters that need to be estimated. To simplify notations, we put the expression in a form where the parameters are arranged in a vector¹¹.

$$\mathbf{v} = \mathbf{K}_{\mathbf{x}} \mathbf{a} \quad \text{where} \quad \mathbf{K}_{\mathbf{x}} = \begin{pmatrix} x & y & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & x & y & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \quad \text{and} \quad \mathbf{a} = \begin{pmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \\ a_5 \\ a_6 \\ 1 \end{pmatrix} \quad (12)$$

In order to do a least square fit, we define a symmetric matrix size 7×7 , as inspired by Farneback¹¹. Using a single 7×7 matrix for least square fit is not a different method from the using an equation systems of size 6×6 ¹². It is merely a different way of arranging the coefficients, that will simplify notations later. We define

$$\mathbf{Q}_{\varphi, \mathbf{x}} = \mathbf{K}_{\mathbf{x}}^T \mathbf{c}_{\varphi, \mathbf{x}} \mathbf{c}_{\varphi, \mathbf{x}}^T \mathbf{K}_{\mathbf{x}}. \quad (13)$$

and then we redefine the error measure in eq. (10) for our parametric motion model,

$$\begin{aligned} \varepsilon(\mathbf{a}) &= \sum_{\varphi=0,45,90,135} \int_{\mathcal{W}} \mathbf{a}^T \mathbf{Q}_{\varphi, \mathbf{x}} \mathbf{a} \, d\mathbf{x} \\ &= \mathbf{a}^T \mathbf{Q}_{tot} \mathbf{a} \end{aligned} \quad (14)$$

where

$$\mathbf{Q}_{tot} = \sum_{\varphi=0,45,90,135} \int_{\mathcal{W}} \mathbf{Q}_{\varphi, \mathbf{x}} \, d\mathbf{x}. \quad (15)$$

The affine parameters that minimize the error measure are computed as

$$\begin{pmatrix} a_1 \\ \vdots \\ a_6 \end{pmatrix} = - \begin{pmatrix} Q_{11} & \dots & Q_{16} \\ \vdots & \ddots & \vdots \\ Q_{61} & \dots & Q_{66} \end{pmatrix}^{-1} \begin{pmatrix} Q_{17} \\ \vdots \\ Q_{67} \end{pmatrix} \quad (16)$$

where Q_{mn} denotes the element at position (m, n) in the matrix \mathbf{Q}_{tot} . As a general result, one might note that this method for affine motions are a paradigm that can be used other parametric motion models. The general case is not described here for readability.

3.2. Handling deformations

We have developed and tested methods to estimate more complicated motions than shift, rotations and deformations. One approach is to use a motion model with more parameters, e.g. quadratic motion model or finite element method. As the number of parameters is increased, the aperture problem gets worse and motion estimates are more susceptible to errors. Therefore it is necessary to either use as few parameters as possible or apply some kind of local-global regularization. We believe in regularization that assumes motions are locally affine, since real world rotations and translations are affine transformations.

3.2.1. Finite Element Motion Model

We have also used the finite element method with linear elements as a parametric motion model. A stiffness matrix provides for regularization. Experimental results are good, but we still have abandoned this approach due to the inferior computational efficiency. Another reason is the difficulty to design stiffness matrices that yield uniform regularization over the entire image, even near the borders of the image. X-ray images may have borders of different shapes, due to the use of collimators. Circular borders are common.

3.2.2. Local-Global Affine Model

Instead of summing the affine 7×7 matrices, \mathbf{Q}_k , over the entire image we compute local sums. Matrices at nearby pixels are given a greater weight than matrices far away. This gives local sums of \mathbf{Q} . In case all weights are positive, the aperture problem is not worse than for the global affine model. The amount of regularization is easily controlled by weights.

4. IMPLEMENTATION

We have implemented the motion compensation algorithm as an integral part of novel angiography software. The angiography software can be started from the ordinary image viewer of SECTRA's commercial PACS*. In order to speed up calculations, the software is prepared to share the workload over four processors, i.e. the number of filter directions. The processing time depends on desired accuracy. For most images, it is adequate to subsample two octaves before applying the motion estimation algorithm. In that case, it takes about a second to compensate an 512×512 image on a dual processor Pentium-II[†] workstation.

*PACS. Picture Archiving and Communicating System.

[†]Pentium-II is a trademark of Intel

4.1. Finding Valid Region

Many X-ray angiography images have been acquired using a collimator. The pixels outside the collimator are dark and lack structure. This region must be considered invalid, otherwise the motion estimation algorithm would track the edge of the collimator rather than the contents of the image. We have applied a simple scheme to find the collimator. It first analyzes image intensity statistics versus the x- and y-axes to find the horizontal and vertical collimators. Then it analyzes the statistics versus the radius to find the circular collimator.

5. EXPERIMENTAL RESULTS

5.1. Experiments on real angiography images

The algorithm has been applied on a large number of image sequences with motion artifacts. Most of these sequences had no deformations and were compensated with such accuracy that no motion artifacts were present in the subtraction images. We have also applied the algorithm on a few sequences where serious deformations are present. In these images, automatic motion compensation gets rid of most of the motion artifacts. We believe that the remaining artifacts are due to structure at different depth are moving in different directions. Figures 3 to 4 shows what happens with our sequence with largest deformations.

A software package with the algorithm has recently been installed at Linköping University Hospital. So far, there has not been any reports that the algorithm has failed for ordinary images of peripheral angiography. We do, however, have a few examples where the algorithm has failed due to extraordinarily large and complicated motions.

5.1.1. Experiment on synthetic images

In clinical use, the contrast injection causes disturbing changes in the image, that may also disturb the motion estimation. We have made an experiment on synthetic images to show that our phase-based motion estimation is less susceptible to such disturbance than the the conventional gradient method, often referred as optical flow⁸. We have used synthetically shifted images to evaluate the accuracy when one of the image frames is disturbed. We have used a popular reference image, Lena256x256, which has been shifted and then subsampled to hide artifacts due to subpixel shifts. The shifts are in all possible directions and we have computed the average performance for all shifts of the same distance.

Since we use iterative refinement, it is most relevant to study performance for shifts less than $\sqrt{2}/2 \approx 0.7$ pixels. After convergence, the warp reduces motion to something less than a half pixel in each of x and y directions.

6. CONCLUSIONS

We conclude that motion automatic motion compensation works for peripheral angiographies. The implications for common practice are that radiologists' time can be significantly reduced in ordinary peripheral angiographies and that the number of retakes due to large or local motion artifacts will be much reduced.

6.1. Diagnostic safety

Several medicals have posed the question whether automatic motion compensation may cause faults in diagnostics. This is still an open issue until we have a long experience of clinical use. We believe that automatic motion compensation with a global affine model does not increase the probability of faults in diagnostics compared to manual pixel shift. It may occur that the automatic motion fails, but that only produces the same characteristic artifacts as a bad pixel shift. A medical will immediately realize it when the motion compensation is bad.

In case we allow large local deformations in the automatic motion compensation, there is an evident risk. E.g. we may make a vessel look much thinner than it is and the medical will see a partial stenosis that does not exist. Therefore, implementations for clinical use should not allow large local deformations.

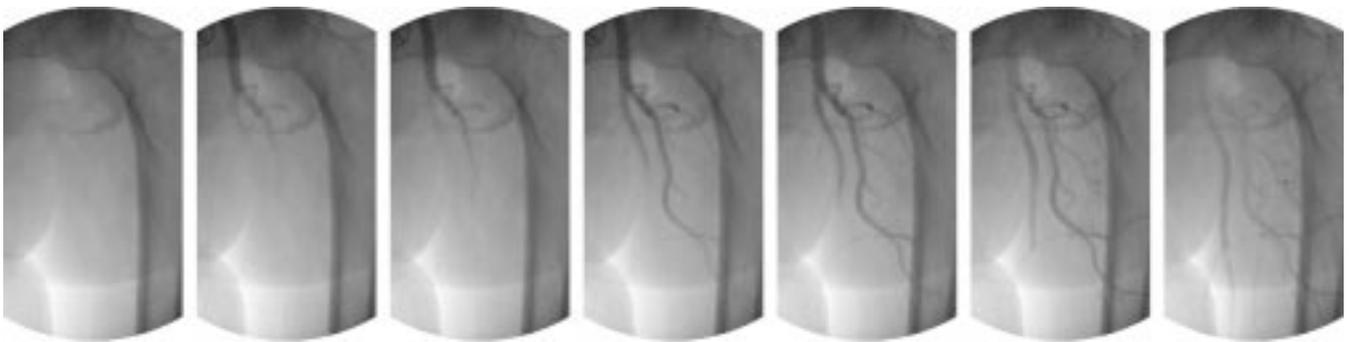


Figure 3. Original X-ray images (severe deformations).

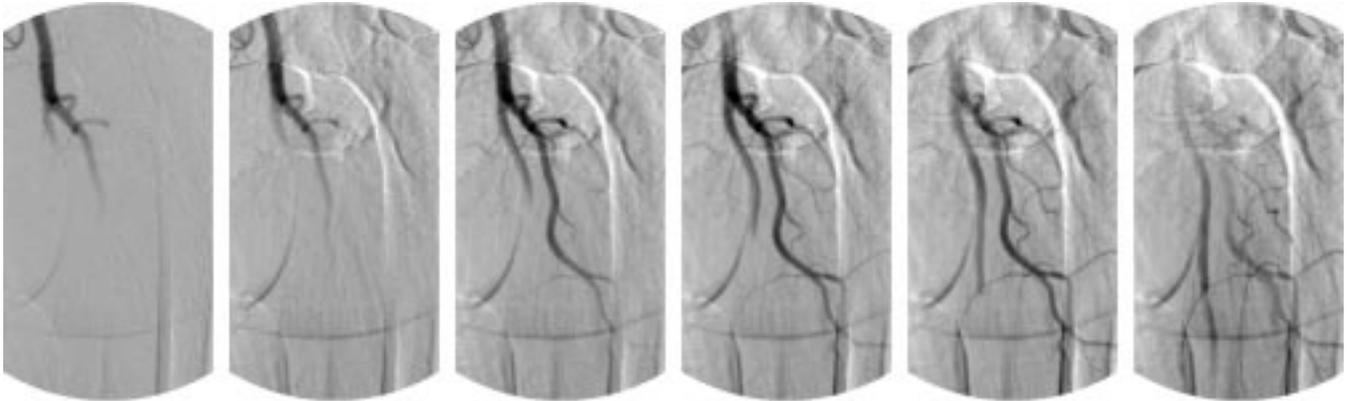


Figure 4. Subtraction Images, no motion compensation

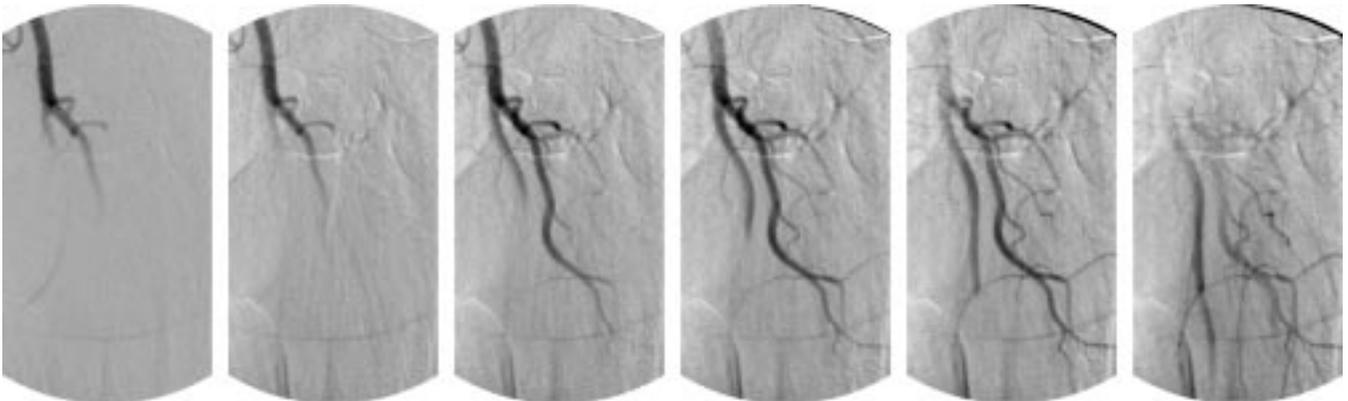


Figure 5. Subtraction after motion compensation. (Remaining artifacts are due to structure at different depth move differently).

6.2. Future research

Because of the parallelism of the algorithm special implementations may for instance allow real time applications, like fluoroscopy. Automatic motion compensation may be helpful in interventional angiography, i.e. angiography where stenoses and other diseases are treated with help from real time X-ray images.

Future research on motion in medicine may enable novel methodology of diagnostics. In theory it seems possible to apply subtraction angiography on a patient in motion, e.g. a walking patient. In the future the basic methodology

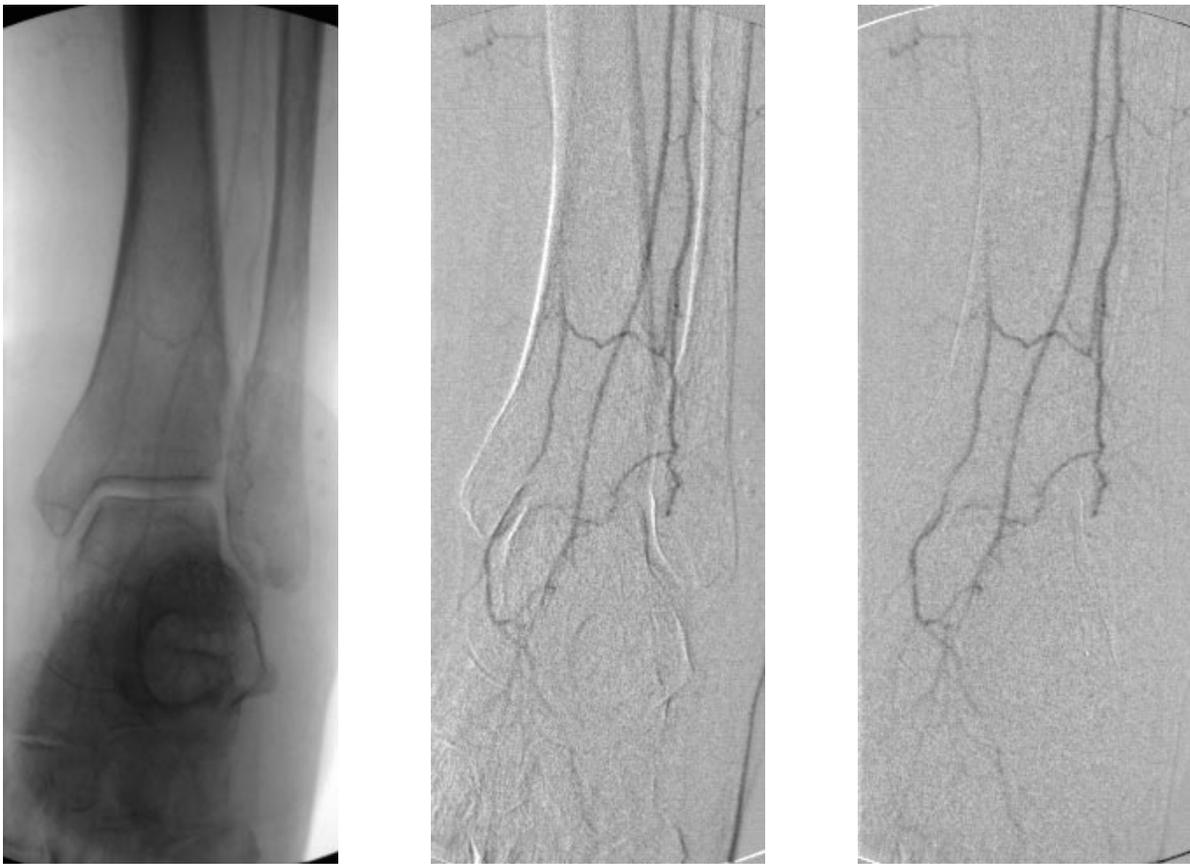


Figure 6. (a) Original image (5/17). (b) Image subtraction without motion compensation. (c) Image subtraction after automatic motion compensation.

can as well be used for estimation of local flow or motion. There is no limitation in our algorithm that restricts the use to two-dimensional images. Motion compensation can as well be applied to three dimensional MR images, and we suggest research on functional MRI and diagnostics of ischemic heart disease.

ACKNOWLEDGMENTS

Financial support is provided by NUTEK (project 9305120-1) and SECTRA.

The following Medicals have been involved: Lars Thorelius, Erik Hellgren (Linköping University Hospital), Torbjörn Andersson, Asgrimur Ragnarsson (Örebro Regional Hospital), Anders Persson, Göran Iwar (Hudiksvall Hospital) Help on practical details have also been received from John Goble and Claes Lundström (SECTRA).

REFERENCES

1. C. f. D. C. Lewis A. Connor, David Satcher and Prevention, "Reducing the burden of cardiovascular disease: Cdc strategies in evolution." Chronic Disease Notes and Reports, 1997.
2. C.-J. Westelius, *Focus of Attention and Gaze Control for Robot Vision*. PhD thesis, Linköping University, Sweden, S-581 83 Linköping, Sweden, 1995. Dissertation No 379, ISBN 91-7871-530-X.
3. D. J. Fleet, A. D. Jepson, and M. R. M. Jenkin, "Phase-based disparity measurement," *CVGIP Image Understanding* **53**, pp. 198-210, March 1991.
4. A. D. Calway, H. Knutsson, and R. Wilson, "Multiresolution frequency domain algorithm for fast image registration," in *Proc. 3rd Int. Conf. on Visual Search*, (Nottingham, UK), August 1992.
5. A. D. Calway, H. Knutsson, and R. Wilson, "Multiresolution estimation of 2-d disparity using a frequency domain approach," in *Proc. British Machine Vision Conf.*, (Leed, UK), September 1992.

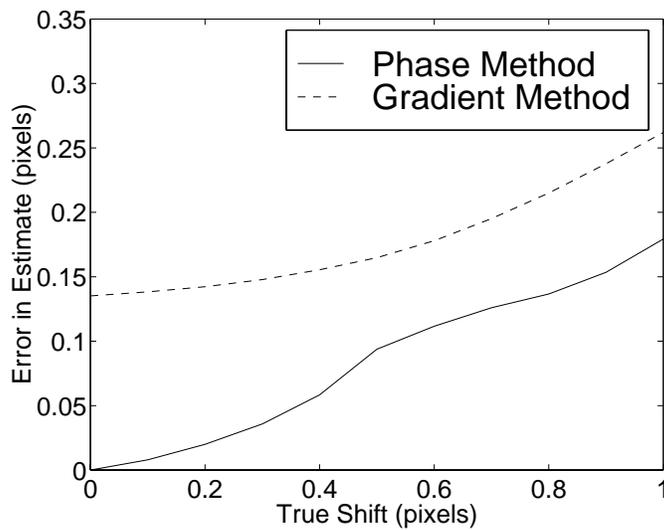


Figure 7. The phase-based method is more accurate than the conventional gradient method. This figure shows a comparison on images (Lena 256x256) that are shifted synthetically (shifting Lena 512x512 before subsampling). One of the image frames has been disturbed by adding a transparent stripe across the image, in order to simulate a contrast bolus. (One pass estimation, i.e. no iterative refinement).

6. D. J. Fleet and A. D. Jepson, "Computation of Component Image Velocity from Local Phase Information," *Int. Journal of Computer Vision* **5**(1), pp. 77–104, 1990.
7. G. H. Granlund and H. Knutsson, *Signal Processing for Computer Vision*, Kluwer Academic Publishers, 1995. ISBN 0-7923-9530-1.
8. B. K. P. Horn and B. G. Schunk, "Determining optical flow," *Artificial Intelligence* **17**, pp. 185–204, 1981.
9. A. D. Jepson and D. J. Fleet, "Scale-space singularities," in *Computer Vision-ECCV90*, O. Faugeras, ed., pp. 50–55, Springer-Verlag, 1990.
10. A. D. Calway, H. Knutsson, and R. Wilson, "Multiresolution estimation of 2-d disparity using a frequency domain approach," in *Proc. British Machine Vision Conf.*, (Leed, UK), September 1992.
11. G. Farneback, "Motion-based Segmentation of Image Sequences," Master's Thesis LiTH-ISY-EX-1596, Computer Vision Laboratory, S-581 83 Linköping, Sweden, May 1996.
12. J. Shi and C. Tomasi, "Good features to track," *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 593–600, 1994.