Image Warping for Accurate Digital Subtraction Angiography

Rajiv Gupta
GE Corporate R&D, 1 River Rd.
K1-5C39, Schenectady, NY 12301
gupta@crd.ge.com

Chukka Srinivas
Imaging Systems, Hewlett Packard Company
3000, Minuteman Road, Andover, MA 01810
srinivas@hpanvl.an.hp.com

Abstract

This paper describes a system for sub-pixel registration of mask and opacified digital X-ray images so that the mask image can be accurately subtracted out from the opacified image. Our technique results in digital subtraction angiograms (DSAs) with fewer artifacts and better contrast.

Introduction

Digital subtraction angiography is a routine X-ray procedure for observing vasculature. In DSA, an X-ray image of normal anatomy is taken before and after an X-ray opaque contrast agent is injected into the blood vessels [Macovski, 1983]. Logarithmic subtraction of the before image from the after image — these images are sometimes referred to as mask and opacified images or pre- and post-contrast images — subtracts out all but the opacified blood vessels. In principle, only the opacified vasculature should be visible in a DSA. However, in practice, because of the following factors, this is rarely the case.

Patient Motion. There is a time lag between the acquisition of pre- and post-contrast images. Small patient motions in this interval cause misregistration between the two images. This leads to motion artifacts in the final DSA.

Involuntary Motion. While certain types of patient motions can be controlled and minimized, several other types of involuntary motions (e.g., cardiac) are much harder to control without major intervention. This fact has restricted or limited the use of DSA in several anatomical regions. For example, for cardiac procedures, only the opacified images are used.

Change in Camera Parameters. Sometimes, it is necessary to move the X-ray imager between pre- and post contrast exposures. Because of inability to bring the imager back to the exact same spot, mechanical vibrations in the structural components of the imager, and several similar effects, the mask and the opacified image are taken under slightly different setting. In general, they should be regarded as images taken by two highly correlated, but different cameras.

Hysteresis. If image intensifiers are used as the sensor, because of the changes in the electrical and magnetic environment the pre- and post-contrast images will be distorted differently. For example, one of our experimental machines uses a rotating gantry fitted with image intensifiers for image acquisition. It gathers mask images in its forward sweep, and opacified images on its reverse sweep. We have observed several artifacts in the DSAs because of the hysteresis effect.

In this paper, a new technique for warping a mask image to accurately register it to the opacified image is described. This work, which is a direct application of research performed in the area of remote-sensing, illustrates the dual-use nature of the basic algorithms involved. The image to image warp, which is used during subtraction, is computed as follows.

1. **Match Point Computation.** In this step, a set of 2-D points in the mask image, and their corresponding points in the opacified image, are derived. The overall procedure used for match point computation is based on the STEREOSYS software developed by Marsha Hanna [Hannah, Oct 1985] and uses the following computation steps.
   - *Derive Image Hierarchy.* The highest resolution mask and opacified images are recursively decimated to half the resolution using Gaussian convolution to derive a pyramid of images.
   - *Compute Interesting Points.* Using a statistical operator, an interest score is assigned to each pixel. The peaks in the interest operator are designated as the interesting points.
   - *Matching.* An attempt is made to match each interesting point to its corresponding point in the opacified image using area-based correlation.
2. **Image to Image Warp Computation.** Using the image to image match points derived above, a 2-D perspective transformation that maps the matched points in the mask image to their corresponding points in the opacified image is computed. This transformation is used to warp the mask image before it is log-subtracted from the opacified image.

Log-Subtraction. From the logarithm of each pixel value in the opacified image, the log of the corresponding pixel value in the transformed (i.e., warped) mask image is subtracted out.

Presently, a 2-D perspective warp is being used d found to be satisfactory for the task at nd [Gupta and Hartley, 1992]. More sophisticated rps such as bi-cubic or thin-plate splines can be :d [Wolberg, 1990]. However, the computation time juirement in the present application ($\leq$ 15 seconds), ;ues for a simple but fast warping algorithm.

:n the following sections we describe the above prosing steps.

Match Point Computation

a preprocessing step, an image hierarchy or pyramid :onstructed in order to accelerate the computation of .tch points. This is accomplished by successively re:ing both images in the stereo pair to half their size ıd resolution) via subsampling using Gaussian convoion. The matching process begins at the bottom of pyramid and works its way up to images with higher l higher resolution. During preprocessing, a set of teresting points" are also computed in one of the ıges. The matching process attempts to match only points in this set with their corresponding points in other image.

n order to match a point in the mask image to its responding point in the opacified image, a small of imagery around it is correlated with all tiles the opacified image. This process proceeds hirchically from the lowest resolution to the highresolution. The center of the tile in the opaci. image which gives the maximum correlation (and ısfies some other confirmation conditions) is de:ed to be the corresponding match point. These processing steps are largely the same as those in EREOSYS testbed and the reader is referred to .nnah, April 1980, Hannah, Oct 1985, Quam, 1987] details.

Image to Image Warp Computation

: input mask and opacified images may be rotated h respect to each other and may be only partially rlapping. We overcome the problems arising because hese effects via a 2D perspective transformation M_I t maps a point $\tilde{u}$ in the first image, to the neighbord of its corresponding match point $\tilde{u}'$ in the second .ge. The following theorem predicts under what conons M_I will completely take out all the distortion in mask image [Hartley *et al.*, 1992].

eorem 1 *Let* $\tilde{u}_i = [u_i, v_i, 1]^T$ *and* $\tilde{u}'_i = [u'_i, v'_i, 1]^T$ *be images of 3-D points* p_i, $i = 1 \ldots n$, *in the given impair. Each* $\tilde{u}_i$ *in the first image can be transformed ts corresponding match point* $\tilde{u}'_i$ *in the second image a 2-dimensional perspective transformation if all* p_i*s in a plane.*

Proof: For all match points $[u_i, v_i, 1]$ and $[u'_i, v'_i, 1]$ which are images of points p_i in a plane, we have to show the existence of a 3×3 matrix $M_I = [m_{ij}]$ such that

$$\begin{bmatrix} w'_i u'_i \\ w'_i v'_i \\ w'_i \end{bmatrix} = \begin{bmatrix} m_{11} & m_{12} & m_{13} \\ m_{21} & m_{22} & m_{23} \\ m_{31} & m_{32} & m_{33} \end{bmatrix} \begin{bmatrix} w_i u_i \\ w_i v_i \\ w_i \end{bmatrix}. \quad (1)$$

Without loss of generality, assume the all p_is lie in the X-Y plane (i.e. the plane $z = 0$) and the camera matrices are denoted by P_1 and P_2. The image coordinates of each $p_i = [x_i, y_i, 0]$, in the two images, are given by

$$[w_i u_i,\ w_i v_i,\ w_i]^T = P_1[x_i,\ y_i,\ 0,\ 1]^T \quad (2)$$

$$[w'_i u'_i,\ w'_i v'_i,\ w'_i]^T = P_2[x_i,\ y_i,\ 0,\ 1]^T \quad (3)$$

Equations (2) and (3)) can be rewritten as

$$\begin{bmatrix} m_{11} & m_{12} & m_{13} \\ m_{21} & m_{22} & m_{23} \\ m_{31} & m_{32} & m_{33} \end{bmatrix} \begin{bmatrix} w_i u_i \\ w_i v_i \\ w_i \end{bmatrix}. \quad (4)$$

$$[w_i u_i,\ w_i v_i,\ w_i]^T = \hat{P}_1[x_i,\ y_i,\ 1]^T \quad (5)$$

$$[w'_i u'_i,\ w'_i v'_i,\ w'_i]^T = \hat{P}_2[x_i,\ y_i,\ 1]^T \quad (6)$$

where $\hat{P}_i$ is the 3×3 matrix formed by deleting the third column of P_i. From these relations it follows that

$$[w'_i u'_i,\ w'_i v'_i,\ w'_i]^T = \hat{P}_2 \hat{P}_1^{-1}[w_i u_i,\ w_i v_i,\ w_i]^T \quad (7)$$

as required. □

In practice, since p_is do not lie in a plane, there would be deviation in the values computed using M_I. Let $\mathcal{P}$ be the plane that fits p_is the best. M_I would account for all disparities arising because of the change in viewpoint except the deviations which are related to a point's distance from $\mathcal{P}$.

The problem of computing M_I can be stated as that of minimizing the sum of errors ϵ_i in the following, possibly overconstrained, system of equations.

$$[u'_i,\ v'_i,\ w'_i]^T = M_I[u_i,\ v_i,\ 1]^T + \epsilon_i. \quad (8)$$

Each match point leads to two equations:

$$m_{11}u_i + m_{12}v_i + m_{13} - u'_i(m_{31}u_i + m_{32}v_i + m_{33}) = 0$$
$$m_{21}u_i + m_{22}v_i + m_{23} - u'_i(m_{31}u_i + m_{32}v_i + m_{33}) = 0$$

This system of equation can be cast as a minimum least-square error solution to $Ax = 0$, where x is a 9 dimensional vector containing the entries of M_I, and A is $2n \times 9$ matrix of known coefficients with n being the number of available tie points. Since the entries of M_I are only determined up to a constant multiplier, the constraint $||x|| = 1$ can be imposed to avoid the all-zero solution.

A rough, initial transformation is first computed based on user-provided tie points between the two images. As few as four tie points are sufficient to start the process. Very often, since the images are well aligned to begin with in the DSA application, the initial transform is taken to be an identity transformation. Using this rough transformation, unconstrained hierarchical matching, as described in [Hannah, April 1980, Hannah, Oct 1985], can proceed.

Once a set of match points has been computed, M_I can be recomputed using computer generated match points. The computation of M_I is rather fast, only requiring solution to the linear system of equations given in Eq. (8). In fact it is possible to refine M_I at any intermediate point during the hierarchical matching process as more match points become available.

Robust warp computation

When hundreds of match points are computed, a few bad matches are inevitable. In order to nullify the effect of these outliers, a statistically robust measure of error is used in the final computation of the image to image warp. Instead of minimizing the mean square error in mapping match points from the mask image to the opacified image, we minimize the *median error*. This is done as follows.

A small sample set of match points is randomly extracted from the full set of match points. Using this sample set — it consists of 5 randomly picked points (4 points are actually sufficient) — an image to image warp is computed as described above. Using this warp, the error in mapping all points from the mask image to their corresponding point in the opacified image is computed. The median error is used as the metric for the goodness of a sample.

If we assume that more than half the match points are actually correct, than the probability that the above sample set of match points consists entirely of good match points is greater than 1 out of 32. Also, if more than 50% of the match points are good, and the image to image warp is computed using a sample consisting entirely of good points, then the median error will be unaffected by the error in the outliers. So the basic procedure is to repeatedly pick a sample set, and compute the median error. The sample that yields the minimum median error after one hundred iterations is used to compute the final image to image warp. Reader can confirm that with high probability, out of one hundred tries, there will be at least one sample which has no outlier.

Experimental Results

Comparison of warped-DSA to the result of direct log-subtraction which assumes that the mask and opacified images are registered with respect to each other shows marked improvement in contrast and reduction in background artifacts. Fig. 1 shows a DSA produced by direct log-subtraction (i.e., assuming that the mask and opacified images are already registered). In this DSA, the opacified image was taken at a time when the opacifying agent had not spread into the blood vessels. So ideally, the subtraction should have yielded a blank picture. However, because of the factors discussed in the introduction, there are several artifacts and boundaries of major organs are clearly visible. Fig. 2 shows a DSA produced by our scheme.

Over twenty DSAs were processed through the system and the following results were observed.

1. **Elastic warping is robust.** In all cases, the matching was robust and the system succeeded in finding an image to image warp.
2. **Elastic warping is superior to no Warping.** The quality of the final DSA was considerably better than the DSA produced when no warping is done (i.e. the case when the mask and opacified images are assumed to be registered with respect to each other and directly subtracted).
3. **Elastic warping is superior rigid displacement.** In many commercial systems, the registration of mask and opacified images is restricted to a global displacement of one image with respect to the other. This global displacement vector is computed iteratively by taking small steps in the direction of increasing correlation between mask and opacified images, and making sure that some consistency conditions are satisfied. Whenever this simple scheme could find a global displacement vector — it does not't converge about 25% of the times — the warped DSA was of superior quality.

References

[Gupta and Hartley, 1992] R. Gupta and R. Hartley. STEREOSYS 3.0 user's guide. Report, GE Corporate R&D, June 1992.

[Hannah, April 1980] M. J. Hannah. Bootstrap stereo. In *Proc. Image Understanding Workshop, College Park, MD*, pages 210–208, April 1980.

[Hannah, Oct 1985] M. J. Hannah. A description of sri's baseline stereo system. Technical Report Tech. Note 365, SRI International Artificial Intelligence Center, Oct. 1985.

[Hartley *et al.*, 1992] R. Hartley, R. Gupta, and T. Chang. Stereo from uncalibrated cameras. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, pages 761–764, 1992.

[Macovski, 1983] A. Macovski. *Medical Imaging System*. Prentice-Hall, Englewood Cliffs, NJ, 1983.

[Quam, 1987] L. H. Quam. Hierarchical warp stereo. In M.A. Fischler and O. Firschein, editors, *Readings in Computer Vision*, pages 80 – 86. Morgan Kaufmann Publishers, Inc., 1987.

[Wolberg, 1990] G. Wolberg. *Digital Image Warping*. IEEE Computer Society Press, Los Alamitos, CA, 1990.

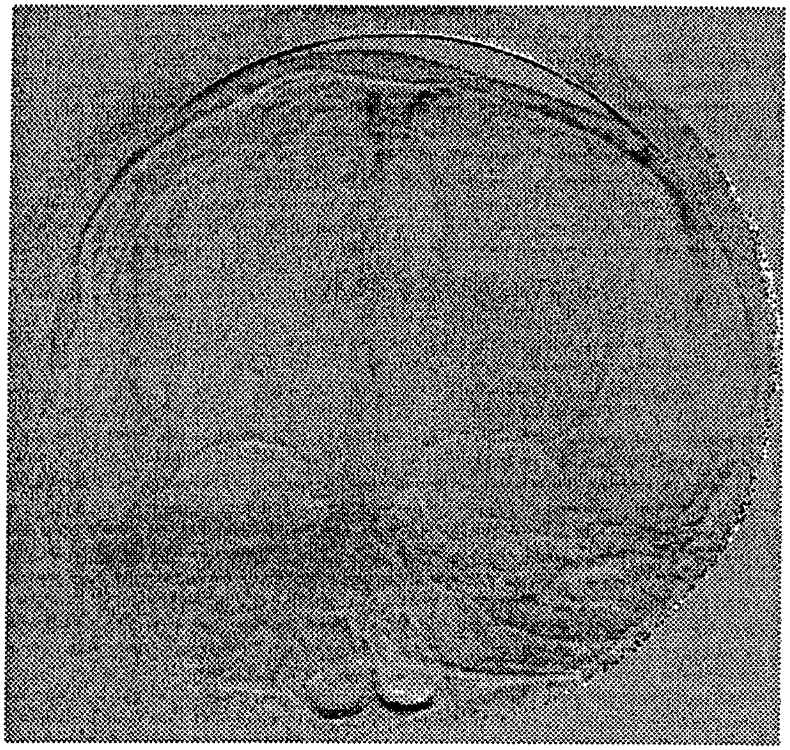

Figure 1: DSA produced by direct subtraction.

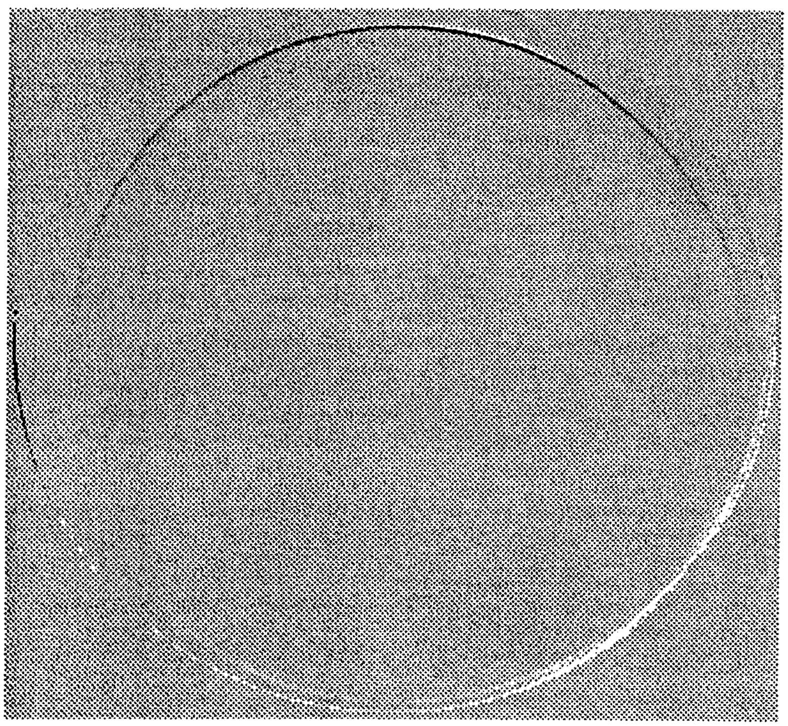

Figure 2: DSA produced by warping the mask image before log-subtraction.