

Image Registration for Digital Subtraction Angiography

Erik H. W. Meijering, Karel J. Zuiderveld, Max A. Viergever

International Journal of Computer Vision, vol. 31, no. 2/3, April 1999, pp. 227–246

Abstract—In clinical practice, digital subtraction angiography (DSA) is a powerful technique for the visualization of blood vessels in the human body. The diagnostic relevance of the images is often reduced by artifacts which arise from the misalignment of successive images in the sequence, due to patient motion. In order to improve the quality of the subtraction images, several registration techniques have been proposed. However, because of the required computation times, it has never led to algorithms that are fast enough so as to be acceptable for integration in clinical applications. In this paper, a new approach to the registration of digital angiographic images is proposed. It involves an edge-based selection of control points for which the displacement is computed by means of template matching, and from which the complete displacement vector field is constructed by means of interpolation. The final warping of the images according to the calculated displacement vector field is performed real-time by graphics hardware. Experimental results with several clinical data sets show that the proposed algorithm is both effective and very fast.

Keywords—Digital subtraction angiography, motion correction, registration, matching, warping.

1 Introduction

In clinical practice, digital subtraction angiography (DSA) is a powerful technique for the visualization of blood vessels in the human body [6, 60]. In ordinary X-ray projection images, blood vessels are hardly visible due to the very low contrast between vessels and the surrounding tissue. This contrast is enhanced by injection of a radiopaque contrast medium (a iodinated solution) into the vessels to be diagnosed. However, without any further processing the contrast between vessels and surrounding tissue is still significantly smaller than that between bone and surrounding tissue. This may introduce severe distortions and, hence, a reduction in the amount of diagnostic information that can be extracted from the images.

In DSA imaging, a *sequence* of images is taken to show the passage of a bolus of injected contrast material through one or more vessels of interest. The contrast distortions in these *live* or *contrast images* are largely removed by subtracting an image taken prior to the arrival of the contrast medium (referred to as the *mask image*). The relatively high image quality, the low patient load as compared to *e.g.* computed tomography angiography (CTA) and the possibility to extract both spatial and temporal information (the dynamic behavior of organs) will likely ensure DSA to remain an important technique.

The subtraction technique is based on the assumption that the tissues surrounding the vessels do not change in position or density during exposure. Clinical evaluations following the introduction of DSA in the early eighties [7, 14, 29] revealed that this assumption is not valid for a substantial number of examinations. Patient motion almost always occurs,

which causes the subtraction images to show artifacts that may hamper proper interpretation of the images. Since then, several techniques have been developed and applied in order to reduce the artifacts and to improve the diagnostic value of DSA images. These include tomographic [6, 34] and tomosynthetic [35] imaging techniques, as well as dual (K -edge) energy subtraction [33, 41], hybrid subtraction [5] and automatic remasking [43]. However, these methods have never been introduced on a large scale in clinical practice, mainly because they require materials and devices which are either very expensive, or difficult to produce, or both.

Alternatively, the artifacts in the subtraction images could be corrected for retrospectively by means of image processing techniques. This is done by calculating the correspondence between pixels in the successive images in the sequence and by warping the images with respect to each other according to this correspondence. This is often referred to as *image registration* and is applied in many other situations where image sequences need to be analyzed. Overviews of registration techniques have been given by Aggerwal & Nandhakumar [1] and Brown [8], and in the field of medical imaging by Van den Elsen *et al.* [55] and Maintz & Viergever [38].

Although many studies have been carried out on this subject over the past fifteen years, they have not led to algorithms which are sufficiently fast so as to be acceptable for integration in clinical applications. In this paper, a new approach to the registration of digital angiographic image sequences is proposed. After a theoretical consideration of the limitations when dealing with *projection* images, in Section 2, the proposed registration approach and some implementation aspects are described in detail in, respectively, Sections 3 and 4, followed by an overview of the complete algorithm in Section 5. Results of experiments on real angiographic image sequences are presented in Section 6 and discussed in Section 7. Concluding remarks are made in Section 8.

2 Theoretical Considerations

Before presenting the method for registration of digital angiographic images, it must be pointed out that these images are in fact *two-dimensional* representations of a *three-dimensional* scene. Because of this, it is often believed that registration techniques that try to recover a complete correspondence (in terms of a displacement vector field $\mathbf{d} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$) between two such images are not likely to be successful [32]. In this section we will investigate the extent to which such a statement holds.

In contrast to *e.g.* range images or images acquired by ordinary photography, in which the grey-values are determined by the (distance to) surfaces of objects in the three-dimensional scene (so called *surface images*), the grey-values in X-ray projection images are determined by the intensity Φ of the X-rays incident on the detector which is, in principle, constituted by the contributions of all particles in the three-dimensional scene. In the case of mono-energetic X-rays, Φ is determined by Lambert-Beer's law:

$$\Phi(x, y) = \Phi_0(x, y)e^{-\mathcal{L}(x, y)}, \quad (1)$$

where $\mathcal{L}(x, y)$ is the line integral $\int \mu(\lambda_{x, y}(\xi))d\xi$ of the linear attenuation coefficient μ along the path traversed by the ray (from the source to the element (x, y) on the detector matrix), of which $\lambda_{x, y} : [0, 1] \rightarrow \mathbb{R}^3$, $[0, 1] \subset \mathbb{R}$ is a parametric representation, and $\Phi_0(x, y)$ is the intensity of the X-rays that is measured when the traversed path is completely in vacuum. In the case of proper calibration with respect to Φ_0 , and logarithmic post-processing of the detected signal, the intensity I of the resulting image at position (x, y) becomes

$$I(x, y) \propto \mathcal{L}(x, y). \quad (2)$$

Given this knowledge, the question arises whether there exists a *two-dimensional* geometrical transformation that completely describes the changes in the images caused by *three-dimensional* object motion. This problem has been studied by Fitzpatrick [18], whose main findings are summarized here. First, the concept of *projected density images* is introduced.

Definition 1 *Let $I(x, y)$ be a two-dimensional grey-value image. If the grey-value at each point (x, y) in the image is proportional to the line-integral projection of the density of a three-dimensional conserved quantity, then $I(x, y)$ is called a projected density image.*

In this definition, a conserved quantity is a quantity whose value within a region can be changed only by transport of particles across the boundary of that region. It is known that the X-ray attenuation coefficient $\mu(x, y, z)$ is proportional to the density ρ of the encountered material in the three-dimensional scene at position (x, y, z) . Then, from (2) it is concluded that Definition 1 applies to X-ray images. The answer to the previously stated question is now provided by the following theorem.

Theorem 1 *Let $I(x, y, t_1)$ and $I(x, y, t_2)$ be two-dimensional projected density images of a three-dimensional scene, acquired at time t_1 and $t_2 > t_1$, respectively, using a device that has a differentiable point-spread function. If the two images are assumed to differ only as a result of the motion of particles in the three-dimensional scene, then there exists a one-to-one two-dimensional mapping $\mathbf{d} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ that transforms points $\mathbf{x}_1 = (x_1, y_1)$ in $I(x, y, t_1)$ to their corresponding points $\mathbf{x}_2 = (x_2, y_2) = \mathbf{d}(x_1, y_1)$ in $I(x, y, t_2)$, and the associated change in intensity (grey-value) is given by the relation*

$$I(x_2, y_2, t_2) = J_r(x_2, y_2)I(x_1, y_1, t_1), \quad (3)$$

where J_r is the Jacobian of the reverse mapping $\mathbf{d}_r : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ that maps points $\mathbf{x}_2 = (x_2, y_2)$ in $I(x, y, t_2)$ back to $\mathbf{x}_1 = (x_1, y_1)$ in $I(x, y, t_1)$.

It should be stressed that the theorem (and its proof [18]) shows only that there *exists* a two-dimensional transformation corresponding to object motion in the original three-dimensional scene. It does not yield any recipe for retrieving this transformation, given the two images. In fact, it is impossible to uniquely retrieve this transformation from the two images only. There are three main reasons for this. (i) Since we are dealing with discrete images, the retrieval of any (analytical) function or its derivatives from the image always requires neighborhood operations. In the case of obtaining a displacement vector in a certain point in one image in a sequence, this implies that if the neighborhood that is involved in the calculation contains several objects that move independently of each other, the result can be expected to be unreliable and inaccurate. Especially in projected density images, these situations are very likely to occur. (ii) At points that lie on isophotes in the image, it is impossible to retrieve the tangential component of the displacement vector. This is generally known as the aperture problem. (iii) The assumption of the two images to differ only as a result of the motion of particles *in* the three-dimensional scene, does not apply to angiography, since in this case new particles (*viz.*, the contrast medium) are introduced in certain parts of the scene. These three problems will limit the accuracy of any registration algorithm for this type of images.

3 Registration Approach

Given a two-dimensional discrete image sequence $I(x, y, t)$ of size $M \times M \times N$, the registration of one of the images, $I(x, y, t_1)$, $t_1 \in [0, N - 1] \subset \mathbb{N}$ with respect to a successive image $I(x, y, t_2)$, $t_1 < t_2 < N - 1$ in the sequence, involves two operations: (i) the computation of the correspondence between the pixels in the two images, and (ii) performing the correction according to this correspondence by warping one of the images with respect to the other. In this section, the proposed approach to perform these operations is presented.

3.1 Control Points Selection

To obtain the correspondence between two images, a possible approach would be to explicitly compute the correspondence for every pixel. However, this is computationally very expensive and will not lead to clinically acceptable algorithms. Even with modern workstations this will still take several hours for images of size 512×512 or 1024×1024 (the usual sizes for angiographic images), not to mention the required computation time for a complete image sequence which usually consists of 15 to 20 of such images.

In order to reduce the computation time to a clinically acceptable level (several seconds), assumptions have to be made concerning the nature of the underlying motion. One possibility would be to assume rigidity of the parts that were imaged and to calculate the correspondence of the two images in terms of only global translation and rotation. Current DSA imaging systems are equipped with a so-called pixel-shifting mode, by which it is possible to correct for gross translational motion only. Although this may locally improve the subtraction in some cases, this will, in general, not yield an overall registration.

The alternative is to assume a certain amount of coherence between neighboring pixels and to calculate the correspondence only for a selected set of so-called *control points*, $\mathbf{p}_i = (x_i, y_i)$, in the image. The overall correspondence can then be obtained by interpolation. The control points can be chosen manually by selection of a region of interest [39,58,64] or can be taken on a regular grid [2,16,53,56,65]. More sophisticated algorithms use image features such as lines [4] or line intersections [51], centers of gravity of closed boundary regions [25], high curvature points (corners) [17], zero crossings of the Laplacian [28], or highly structured regions [50,54].

Since in DSA images artifacts appear only in those regions where strong object edges are present in the individual images in the case of a misalignment between mask and live image, and because edges can be matched better than more homogeneous regions, the selection of control points should be based on an edge-detection scheme. By selecting control points from important edges in the image, the implicit assumption is made that the displacement of points in between edges can be described by the displacements of points *on* these edges. Physically, this means that the displacement of tissue in between objects with a very different density (usually bone/soft-tissue transitions), is dictated by the displacement at these transitions, which indeed seems a valid assumption. Compared to the rubber sheet masking approach of Zuiderveld *et al.* [65], in which control points were taken on a regular grid, this approach has three major advantages: (i) the control points are chosen at those positions where the artifacts can be expected to be the largest, (ii) since the neighborhoods of those points are known to be structured, the reliability of the displacement estimates will be higher, and (iii) the number of control points is reduced, thereby reducing the required computation time.

The location of edges can easily be computed by detecting local maxima of the grey-level gradient magnitude. It should be noted that edges are scale-dependent image features

and therefore can be detected properly only by using derivatives calculated at the proper scale. The commonly used scale-dependent detector for this purpose is the derivative of the Gaussian [12, 27, 40]. Using this approach, the two-dimensional (regularized) gradient magnitude is calculated as

$$\|\nabla L(x, y)\| = \sqrt{(\partial_x L(x, y))^2 + (\partial_y L(x, y))^2}, \quad (4)$$

with

$$\partial_i L(x, y) = (I * \partial_i G)(x, y), \quad (5)$$

where ∂_i denotes the first derivative in the i -th direction and G the bivariate Gaussian.

Note that this edge-detection approach cannot be applied to the contrast images, since in that case the boundaries of vessels would also be detected. The same argument holds for the application of edge detection to the subtraction images. In addition, in the latter case, regions that are already registered but in which there are very strong edges in the underlying mask and live images, will not be identified as potential problem regions. This means that in these regions the interpolation operation may even *introduce* motion artifacts. Therefore, the gradient-magnitude computations should be applied to the mask image. This has the additional advantage that the detection of regions that may give rise to motion artifacts in the subtraction image needs to be carried out only once (as a preprocessing step), instead of repeatedly for every new contrast image.

In principle, the severeness of a motion artifact is directly related to the strength of the underlying edge (*i.e.*, the gradient magnitude $\|\nabla L\|$). In order to be able to extract adequate control points from the gradient-magnitude image, it is required to indicate which edges are important enough to be considered further. This is done by thresholding the gradient magnitude at a value Θ_e , resulting in a binary image in which there are complete regions (contiguous pixels) that are of interest. It is obvious that it is neither manageable (because of the unacceptably large number of points involved) nor necessary (because of the assumed coherence between neighboring pixels) to take all of these edge points as control points for the construction of the warp transformation. Therefore, an additional selection mechanism is required, which will be described hereafter. In order to preserve a sufficient amount of detail prior to this final selection, the scale σ at which the gradient magnitude is calculated should be taken small.¹

Under the assumption of coherence between neighboring pixels in an image, the control points can be constrained to have a minimum distance with respect to each other, say D_{\min} . This minimum distance is related to the image size² according to $D_{\min} = \phi_{\min} M$, in which the minimum distance factor ϕ_{\min} is a parameter of the algorithm. In order to avoid the complete absence of control points in very large regions in which no important edges are detected, the control points should also be constrained to have a maximum possible distance, $D_{\max} = \phi_{\max} M$.

In usual angiographic images, certain regions in the image (at the borders) are of constant grey-value. This is caused by the physical properties of the acquisition system: the X-rays are detected by an image intensifier which has a circular shape. That part

¹It should be noted that taking $\sigma = \sigma_0$ for $M = 512$, implies that σ must be $2\sigma_0$ for $M = 1024$ in order to have the same amount of regularization. However, even when σ is taken twice as large for $M = 1024$, the actual grey-values in the resulting gradient-magnitude image will still be twice as small with respect to those of a 512×512 image of the exact same scene. This should be accounted for in Θ_e .

²The reason for relating D_{\min} and D_{\max} to the image size is that the distance (expressed in pixels) between coherent structures in an image of size 1024×1024 pixels is indeed twice as large as the corresponding distance in an image of size 512×512 pixels of the same scene.

of the image in which the grey-values are due to exposure of X-rays will be called the *exposure region*, \mathcal{R}_E , in the sequel:

$$\begin{aligned} \mathcal{R}_E = \{ \mathbf{x} \mid & (x - x_C)^2 + (y - y_C)^2 \leq R^2 \wedge \\ & x \geq X_{\min} \wedge x \leq X_{\max} \wedge \\ & y \geq Y_{\min} \wedge y \leq Y_{\max} \}, \end{aligned} \quad (6)$$

where (x_C, y_C) is the center of the image intensifier of the DSA system which, in the case of proper calibration, is also the center of the image, R is the radius of the image intensifier and X_{\min} , X_{\max} , Y_{\min} and Y_{\max} are the left, right, upper and lower borders respectively.³ Since the parts of the image outside \mathcal{R}_E do not contain any information that could be used in the registration, the control points should be positioned inside this exposure region, at a minimum distance $D_{\text{exp}} = \phi_{\text{exp}}M$ from the border $\partial\mathcal{R}_E$ of the region. An example of control point selection based on this approach is shown in Figure 1. More details will be given in Section 4.

3.2 Displacement Calculation

Techniques for the calculation of the displacement (or motion) of certain structures in one image with respect to another can be divided into two categories: (i) gradient-based *optic-flow* techniques, and (ii) *template-matching* based techniques. In this section, both techniques will be discussed with respect to their applicability to the computation of displacements in digital angiographic images.

3.2.1 Optic-Flow Based Approach

Optic-flow techniques are explicitly based on the assumption that motion of objects in the images cause a change only in *position* of the corresponding grey-level patterns, while the patterns themselves remain unchanged. For a two-dimensional image sequence $I(x, y, t)$ this implies that $I(x, y, t) = I(x + \delta x, y + \delta y, t + \delta t)$, where δx , δy and δt denote small changes in position and time, respectively. Expanding the right-hand side of this equation, dividing by δt and taking the limit $\delta t \rightarrow 0$, we arrive at:

$$\frac{dI}{dt} = \frac{\partial I}{\partial t} + \nabla I \cdot \mathbf{u} = 0 \quad (7)$$

where ∇I denotes the two-dimensional gradient of the grey-level image I , and $\mathbf{u} = (u, v)$ denotes the velocity in the image plane. Equation (7) is generally known as the *optic-flow constraint equation* [30].

It is well known that the optic-flow problem defined in this way is ill-posed, by the fact that the solution to this problem is not unique (one equation in two unknown variables

³The grey-values outside \mathcal{R}_E are of a constant value which is often much larger than those inside the region. This makes it very easy to extract the five parameters, by evaluating scan-lines from the border to the center (x_C, y_C) of the image. For example, in order to extract the radius R , scanning starts at pixel $(0, 0)$. During scanning, the grey-level at the next position along the scan-line is compared to the value at the previous position. As soon as a difference is detected, say at position (x_i, y_i) , scanning is terminated and the radius is computed as $R = \sqrt{(x_C - x_i)^2 + (y_C - y_i)^2}$. (Of course, it would also be possible to scan in the reverse direction (from the center to the border), which will immediately yield R . However, this would be computationally more expensive since in usual X-ray images the path length from (x_C, y_C) to (x_i, y_i) is far larger than that from $(0, 0)$ to (x_i, y_i) , and the additional scanning time does not outweigh the time required to compute the squares and the square root.) The remaining four parameters X_{\min} , X_{\max} , Y_{\min} and Y_{\max} are found in a similar way, using the starting points: $(0, \frac{1}{2}M)$, $(M, \frac{1}{2}M)$, $(\frac{1}{2}M, 0)$ and $(\frac{1}{2}M, M)$, respectively.




Figure 1. An example of control points selection. Top-left: the mask image $I(x, y, 0)$ of a $1024 \times 1024 \times 18$ digital angiographic image sequence (data set CER in Table 1). Top-right: the subtraction of the mask and one live images under contrast enhancement, showing motion artifacts. Bottom-left: the thresholded gradient-magnitude image of the mask image as a prediction of the location of motion artifacts in the subtraction image (white pixels correspond to pixels whose grey-value is above the threshold in the original grey-value gradient-magnitude image). Bottom-right: the set of selected control points P (white dots), with the thresholded gradient-magnitude image of the mask (now shown in 50% grey) superimposed.

u and v). Only the component of \mathbf{u} in the gradient direction ∇I can be calculated since a displacement in the tangential direction will never lead to a change in the grey-level distribution. This is often referred to as the aperture problem. In order to be able to obtain a unique solution for u and v , additional constraints need to be imposed, *e.g.* by requiring the displacements (or velocities) $u(x, y)$ and $v(x, y)$ to vary smoothly within the image. Several approaches have been described in the literature [21, 28, 30].

However, the just described optic-flow technique cannot simply be applied to X-ray projection images. As mentioned in Section 2, these images are two-dimensional projections of three-dimensional scenes. A thorough analysis by Fitzpatrick [18] has indicated that for these images, the total time-derivative is

$$\frac{dI}{dt} = -I\nabla \cdot \bar{\mathbf{u}}, \quad (8)$$

where $\nabla \cdot$ denotes the divergence operator and $\bar{\mathbf{u}}$ the weighted average (along the rays), of the velocities perpendicular to the rays. This contradicts the assumption made in (7). In addition, it has to be noted that in the particular case of *angiographic* image sequences, the optic-flow approach suffers from all of the three problems mentioned at the end of Section 2. Taking into account these considerations it can be expected that the optic-flow technique will not yield accurate results when applied to these images. This was confirmed by pilot experiments carried out in our group.

3.2.2 Template-Matching Based Approach

Template-matching techniques are based on the assumption that the displacement \mathbf{d} of every pixel in an image $I(x, y, t_1)$ can be approximated by taking a small window (or neighborhood) \mathcal{W} of $W \times W$ pixels around the pixel and finding the corresponding window in a successive image $I(x, y, t_2)$, $t_2 > t_1$ in the sequence by optimization of the measure of match (according to a certain similarity criterion) under translation. In principle, these approaches also suffer from the first two problems mentioned in Section 2 (*viz.*, independently moving structures and the aperture problem). However, they can be made more robust against the inflow of contrast than optic-flow techniques, by a proper choice for the similarity criterion.

Similarity Criterion A crucial aspect of any template-matching method is the similarity criterion that is used to determine the amount of correspondence between regions in successive frames. Many criteria have been devised for this purpose, including normalized cross-correlation [3, 64], the correlation coefficient [3, 45], statistical correlation [47], stochastic and deterministic sign change [31, 57, 65], the sum (or mean) of absolute valued differences [19, 39, 56], coincident bit counting [13, 59] and the variance of (squared) differences [15, 16]. However, most of these criteria are based on the same assumption as optic-flow techniques, *viz.*, that the structures in the image change only in position and not in intensity (*i.e.*, $\nabla \cdot \bar{\mathbf{u}}$ is zero). As mentioned before, this is not a valid assumption when dealing with angiographic image sequences. In addition to that, these criteria are quite sensitive to the inflow of contrast.

Criteria that are very robust against this phenomenon are those based on the histogram of differences [10, 11]. With these criteria, advantage is taken of the fact that in the case of optimal alignment, only a small number of difference grey-levels have a high relative frequency while the majority of grey-levels have a low relative frequency, which results in a sharply peaked histogram. This is true whether or not the window \mathcal{W} contains opacified vessels, the former case resulting in two peaks and the latter in only one peak. In the case of misalignment, the histogram will have a larger dispersion in both cases (Figure 2).

For every pixel in an image and for every displacement \mathbf{d} of a window \mathcal{W} around that pixel, the amount of dispersion of the histogram $\mathcal{H}(\delta)$ of differences δ in \mathcal{W} can be computed by a weighted summation over the bins (which are assumed to have size 1 here,




Figure 2. Normalized histograms for two different regions of interest (indicated by the white boxes, 75×75 pixels in size). Left column: a region without (top) and with (bottom) contrasted vessels. Middle column: the corresponding normalized histograms in the case of alignment of mask and contrast image. Right column: the corresponding normalized histograms in the case of misalignment of mask and contrast image.

i.e., $\delta \in \mathbb{Z}$ and $\mathcal{H} : \mathbb{Z} \rightarrow \mathbb{R}_0^+$), which is expressed in the measure

$$\mathcal{E}(\mathbf{d}) = \sum_{\delta=\delta_{\min}}^{\delta_{\max}} f(\mathcal{H}(\delta)), \quad (9)$$

where $[\delta_{\min}, \delta_{\max}] \subset \mathbb{Z}$ is the range of all possible difference values, $\mathcal{H}(\delta)$ is the fraction (or relative frequency) of pixels in the window \mathcal{W} that have a difference value of δ , $f : \mathbb{R} \rightarrow \mathbb{R}$ is the weight function and $\sum_{\delta=\delta_{\min}}^{\delta_{\max}} \mathcal{H}(\delta) = 1$ because the histogram must be normalized in order to have a fair comparison between the values of this measure for different displacements \mathbf{d} . The constraints that should be imposed on the weight function f are given by the following axiom.

Axiom 1 *Let $\mathcal{H}(\delta)$ be a normalized histogram, *i.e.*, $\sum_{\delta} \mathcal{H}(\delta) = 1$, of all possible difference values $\delta \in [\delta_{\min}, \delta_{\max}] \subset \mathbb{Z}$ in a window \mathcal{W} around a certain pixel in an image, and let $\mathcal{E}(\mathbf{d}) = \sum_{\delta} f(\mathcal{H}(\delta))$ be a measure expressing the amount of dispersion of the histogram for a displacement of \mathbf{d} . In order for $\mathcal{E}(\mathbf{d})$ to be an adequate similarity measure for the purpose of registration, the following constraints must be imposed:*

1. *the weight function f must be such that the measure \mathcal{E} assumes its minimum value if and only if for all $\delta \in [\delta_{\min}, \delta_{\max}]$ we have $\mathcal{H}(\delta) = (\delta_{\max} - \delta_{\min} + 1)^{-1}$,*
2. *the weight function f must be generalized super-additive, which means that for all $a_i \in \mathbb{R}$, $i = 1, 2, 3$ with $0 < a_1 \leq a_2 < a_3$ we have*

$$f(a_2 - a_1) + f(a_3 + a_1) > f(a_2) + f(a_3). \quad (10)$$

When the difference values in the window \mathcal{W} behave like white noise (*i.e.*, the grey-values in the successive images in the sequence are completely uncorrelated), the similarity measure

should assume its smallest value. This is expressed in the first constraint. The second constraint is necessary in order to let the measure of match increase when the histogram becomes more clustered. As proven by Buzug *et al.* [11], these requirements are met by the class of differentiable, strictly convex functions, which are defined as follows.⁴

Definition 2 *A function $f : \mathbb{R} \rightarrow \mathbb{R}$ is called strictly convex if for all $x, y \in \mathbb{R}$, $x \neq y$ and for all $\alpha \in (0, 1)$, the following relation holds:*

$$f(\alpha x + (1 - \alpha)y) < \alpha f(x) + (1 - \alpha)f(y). \quad (11)$$

Subpixel Accuracy The just described matching approach determines displacements only up to integer accuracy. However, as indicated by a clinical evaluation in the early 1980s [7], subpixel misalignments often produce significant artifacts in the subtraction image. To obtain subpixel accuracy, an obvious approach would be to use the measures of match at integer displacements in an interpolation scheme in order to construct a continuous bivariate function $\mathcal{M}(\mathbf{d})$ (referred to as the match surface), and to calculate the global maximum of this function analytically. Of course, bilinear interpolation cannot be used for this purpose since a bilinearly interpolated surface still has its maximum at the integer maximum position. For higher order surfaces, the position of this extreme must be determined by solving $\nabla \mathcal{M}(\mathbf{d}) = \mathbf{0}$. However, it can easily be seen that this will lead to the problem of solving an algebraic equation of degree $n \geq 5$ (in either x or y and with symbolic coefficients that are functions of the measures of match at integer displacements), which cannot be done. Therefore, several authors [16, 58, 65] simplified the problem into constructing two separate mono-variate functions $\mathcal{M}(d_x)$ and solving $\partial_x \mathcal{M}(d_x) = 0$ and $\partial_y \mathcal{M}(d_y) = 0$ to obtain the x and y coordinates of the maximum.

Although these analytical methods lead to an estimation of the displacement with subpixel accuracy, they differ fundamentally from the approaches in which the measures of match for subpixel displacements are calculated explicitly, such as the algorithms described by *e.g.* Yanagisawa *et al.* [64] and Van Tran & Sklansky [56]. In that case, one of the two images needs to be reconstructed and the calculations are carried out using a resampled version of this image, where the resampling is done on a corresponding subpixel displaced grid. Yanagisawa *et al.* [64] and Van Tran & Sklansky [56] reported that an accuracy of 0.1 pixel is sufficient for angiographic images. From our experiments it was concluded that in all cases this image-interpolation approach (even with an accuracy of no more than 0.1 pixel) leads to better registrations compared to the analytical match-interpolation methods (see Figure 3 for examples).

3.3 Displacement Interpolation

In order to be able to carry out the warping of the mask image with respect to the contrast image, it is required to have a complete description of the displacement vector field $\mathbf{d} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$, *i.e.*, the displacement \mathbf{d} must be known for every point \mathbf{p} in the image. So far, the displacements have been calculated only for a selected amount of control points, \mathbf{p}_i ,

⁴It should be mentioned that in Axiom 1, the first constraint, *viz.*, that \mathcal{E} must assume its minimum value in the case of a uniform histogram, may be replaced by the requirement that \mathcal{E} must assume its maximum value in that case. This can only be done if at the same time the requirement of super-additivity is replaced by sub-additivity (which is obtained by replacing the $<$ sign by $>$ in (10)). These new requirements are met by the class of differentiable, strictly *concave* functions (which are obtained by replacing $<$ by $>$ in (11)). In this case, \mathcal{E} should be referred to as a measure of *mismatch* since it assume its maximum value when the windows are most dissimilar.




Figure 3. Two examples of registrations with subpixel accuracy in a critical region of interest, using both the image-interpolation and match-interpolation method. Left column: the original subtractions of unregistered digital angiograms. Middle column: the subtractions after registration using the match-interpolation method (quadratic interpolation). Right column: the subtractions after registration using the image-interpolation method (bilinear interpolation) with an accuracy of 0.1 pixel.

under the assumption that the remainder of the field could be obtained by interpolation. In order to reduce the required computation time for this operation to a minimum, the parameters ϕ_{\min} , ϕ_{\max} and ϕ_{\exp} , as introduced in Section 3.1, should be chosen in such a way that it is allowed to use *linear* interpolation.

Given the set $P = \{\mathbf{p}_i\}$ of control points as extracted by the feature detection and selection algorithm (Section 3.1), a suitable tessellation is required in order to be able to carry out the linear interpolation of the displacements in between these control points. In the case of a regular grid of data points, quadrilaterals are the commonly used polygons [16,20,56,65]. However, in the case of irregularly distributed (or scattered) control points, such a mesh is not guaranteed to exist. The only possible polygons that can be used for this purpose are triangles, so that the control points become the vertices of an irregular triangular mesh. Although triangular meshes are well-known geometric constructions and have been studied and applied by several authors in image registration [23,24], warping and morphing [48,49,63], computer graphics and scientific visualization [22], they have, to our knowledge, never been applied to digital angiographic image sequences.

In constructing the triangulation of a set of data points it should be noted that the solution to this problem is not unique, as can easily be observed from the fact that any quadruple of vertices can always be triangulated in two ways (there are two possible diagonals in the quadrilateral). In order to obtain a unique triangulation that is consistent




Figure 4. An example of mesh generation. Left: the Delaunay triangulation $\mathcal{D}(P)$ of the set P of control points as shown in Figure 1, with the thresholded gradient-magnitude image of the mask (shown in 50% grey) superimposed. Right: the same triangulation superimposed on the original mask image.

with the choice for an interpolation scheme, additional constraints need to be imposed. As stated by Watson & Philip [62], triangles with highly acute interior vertex angles should be avoided in the tessellation since the vertices of these elongated triangles are not capable of reflecting the local variation of the interior points in the dependent variables.

A suitable tessellation for this purpose is the so-called Delaunay triangulation $\mathcal{D}(P)$. It guarantees the smallest of the three angles of a triangle to be as large as possible [36,37] and is unique, except in the degenerate case where four or more points are co-circular [61]. In that case, the Delaunay triangulation is locally non-unique. In our implementation, we used the incremental algorithm described by Watson [61]. An example of the Delaunay triangulation of a set of control points is presented in Figure 4. Additional remarks will be made in Section 4.2.

3.4 Inter-Image Displacement Prediction

So far, the discussion has mainly been focused on the registration of only *two* images $I(x, y, t_1)$ and $I(x, y, t_2)$, $t_2 > t_1$ with respect to each other. However, as indicated at the beginning of Section 3, digital angiographic data sets are complete discrete image sequences of size $M \times M \times N$, where N is usually 10 to 20. The successive images in the sequences are highly correlated since they are projections of the same scene.

Under the assumption that the velocity of patient motion is small with respect to the frequency at which the images are recorded (*i.e.*, the changes from one image $I(x, y, n)$, $n \in [0, N-2] \subset \mathbb{N}$ to a successive image $I(x, y, n+1)$ are small with respect to the total motion from $I(x, y, 0)$ to $I(x, y, N-1)$), it is reasonable to assume that the displacement vector fields $\mathbf{d}(\mathbf{p}, n)$ and $\mathbf{d}(\mathbf{p}, n+1)$ of successive images are highly correlated as well. This can be taken advantage of when computing the correspondence of all images in the sequence with respect to a single image (the mask), or the correspondence of all images with respect to their predecessor in the sequence. In our implementation, the displacements $\mathbf{d}(\mathbf{p}_i, n)$

were used as an estimate for the calculation of $\mathbf{d}(\mathbf{p}_i, n+1)$. More (implementation related) details are provided in Section 4.4.

4 Implementation Aspects

In the previous section, the several components of the registration algorithm have been presented. In this section, several aspects concerning the implementation of these components will be discussed in more detail.

4.1 Control Points Selection

In Section 3.1 it was proposed to calculate the gradient magnitude of the mask image of the sequence as a prediction of the regions in the subtraction image where artifacts can be expected to appear in the case of a misalignment. When dealing with 1024×1024 images, the scale at which the gradient magnitude is calculated should be twice as large as for 512×512 images, in order to have the same amount of regularization. This means that, for 1024×1024 images, the required computation time is at least a factor 8 larger (a factor 4 from the image size and another factor 2 from the kernel size, not to mention the possibility of worse caching behavior at these image sizes).

In order to reduce the required computation time, the gradient magnitude can be calculated using a subsampled version of the mask image, at the cost of loss of accuracy in the positioning of the control points in the final selection. In our implementation, the edge detection is always carried out on an $M_{\text{edge}} \times M_{\text{edge}}$ sized mask image, regardless of the original size. This has the two advantages: (i) regardless of original images size, the computation time is fixed, (ii) the threshold Θ_e can be the same for both image sizes.

The final selection of control points from the gradient magnitude of the mask image is carried out as follows. The image is divided into blocks of size $D_{\text{max}} \times D_{\text{max}}$. In turn, these blocks are sub-divided into smaller blocks of size $D_{\text{min}} \times D_{\text{min}}$. For every large block (from the upper left to the lower right of the image), every small block (again, from upper left to lower right within the large block) is scanned for pixels with a gradient-magnitude value above the specified threshold Θ_e . From these pixels, the one with the largest value is taken as a candidate control point. (If no edge pixels are encountered, no candidate is selected.) The candidate becomes a control point if it is positioned inside the exposure region as defined in (6) and at a distance of at least D_{exp} from the border $\partial\mathcal{R}_E$ of that region. In order to enforce a minimum distance between control points, the gradient-magnitude values in a $(2D_{\text{min}} + 1) \times (2D_{\text{min}} + 1)$ region around the selected point (that point being the center) are suppressed. If no point is selected after a large block has been scanned, the point with the largest gradient-magnitude value in a small block around the center of the large block is taken as a control point so as to constrain the maximum distance between the points.

4.2 Control Points Triangulation

In the eventual warping of the image, the displacements $\mathbf{d}(\mathbf{p}_i)$ of the selected control points \mathbf{p}_i are linearly interpolated to construct the complete displacement vector field $\mathbf{d}(\mathbf{p})$. To this end, it is required that the triangulation $\mathcal{D}(P)$ completely covers the image. However, this is not the case if P consists only of control points from inside the exposure region. In our implementation this is solved by means of four additional corner points that lie outside the image. The displacement of each of these points is explicitly set to zero.




Figure 5. The energy measure \mathcal{E} as a function of the displacement $\mathbf{d} = (d_x, d_y)$ in x and y direction, using different window sizes. Left: result for a window size of 11×11 pixels (43 local extrema). Middle: result for a window size of 31×31 pixels (although hardly visible there are still 6 local extrema). Right: result for a window size of 51×51 pixels (one local extreme, corresponding to the correct displacement). As can be seen from this figure, the window size determines the smoothness of the function $\mathcal{E}(\mathbf{d})$ and thereby the number of local extrema in which the optimization technique might become trapped.

As will be explained in Section 4.5, special attention needs to be paid to the warping of triangles (referred to as corner triangles) of which at least one of the vertices is a corner vertex. In order to remove serious border artifacts that result from the warping, a set of equally spaced points that are positioned exactly on the exposure region border $\partial\mathcal{R}_E$ are taken as additional control points for which the displacement will be calculated.

4.3 Similarity Measure Computation

As proposed in Section 3.2.2, the displacement of every control point is calculated by means of template matching, using the histogram-of-differences similarity measure of Buzug *et al.* [11]. Since the largest part of the computation time is due to the calculation of this measure, a few notes should be made concerning its implementation.

First, the time required to compute the measure of match $\mathcal{E}(\mathbf{d})$, as defined in (9), should be as small as possible, which implies that the function $f(\mathcal{H}(\delta))$ should not only satisfy the requirements of differentiability and convexness, but should also be computationally cheap. To this end, the energy function $f(\mathcal{H}(\delta)) = \mathcal{H}^2(\delta)$ was chosen, since it involves only a single multiplication [9, 11].

A very important parameter is the size of the window \mathcal{W} . In order to reduce the computation time to a minimum, this window should be as small as possible. However, in principle, the size of the window determines the amount of statistical information that is provided and, therefore, the smoothness of the function $\mathcal{E}(\mathbf{d})$. In order to allow for computationally cheap optimization techniques, such as hill-climbing (or the method of Powell [46]), it is a prerequisite that the function $\mathcal{E}(\mathbf{d})$ is sufficiently smooth so as to have only *one* extreme, corresponding to the actual displacement. Small windows will yield unreliable match values, which causes the resulting function $\mathcal{E}(\mathbf{d})$ to be rather coarse (many local extrema). This is illustrated in Figure 5. Experiments for several thousands of points have indicated that a window size of 51×51 pixels yields a good compromise between computational speed and statistical reliability.

Furthermore, digital angiographic images usually have a grey-value resolution of 10 bits, *i.e.*, $I(x, y, t) \in [0, 1023] \subset \mathbb{N}$. This means that the difference values are in the range $[-1023, 1023]$. In order to avoid time-consuming checks in filling the histogram, the histogram should have the same range, with a bin-size of 1 (no clustering of bins).

The computation time can be reduced somewhat further by reducing the summation range $[\delta_{\min}, \delta_{\max}]$. As can be seen from Figure 2, the majority of difference values is concentrated in the range $[-200, 100]$. Restricting the summation to this range reduces the computation time without affecting the accuracy of the calculation.

Finally, in order to let the displacement computations be based on actual image information only (*i.e.*, grey-values inside the exposure region), the difference values outside the exposure region are not incorporated in the computations.

4.4 Inter-Image Displacement Prediction

In Section 3.4 it was stated that when successively computing the displacement vector field $\mathbf{d}(x, y, n)$ for all images $I(x, y, n)$, $n \in [1, N - 1] \subset \mathbb{N}$ with respect to a single image $I(x, y, 0)$ (the mask), the displacements obtained in the previous iteration, $\mathbf{d}(x, y, n - 1)$, can be used as an estimate for the displacements in the current image. If the assumptions concerning the patient motion hold, the use of these estimates may reduce the number of iterations of the hill-climbing algorithm in finding the optimal correspondence.

As proposed in Section 3.2.2, the computation of displacements is carried out hierarchically, *i.e.*, first up to an integer number and then with subpixel accuracy. However, the computation of the fractional part of the displacements requires interpolation of one of the images (bilinear interpolation was used in our implementation). This is a relatively expensive operation. To avoid the need to interpolate the image when calculating the integer part of the displacement, only that part was used in the predictions.

4.5 Image Warping

Because of the coherence in behavior of neighboring pixels, it was proposed in Section 3.1 to calculate the displacement between two images from the sequence only for a selected set of control points \mathbf{p}_i . In Section 3.3 it was advocated that in order to reduce the computation time to a minimum, this should be done in such a way that it is allowed to use linear interpolation to obtain the displacement $\mathbf{d}(\mathbf{p})$ for any other point \mathbf{p} in the image, and an approach to realize this was presented.

The actual warping of an image $I(x, y)$ is carried out as follows. For every triangle Δ_{ijk} in the mesh, the constituent vertices \mathbf{p}_i , \mathbf{p}_j , and \mathbf{p}_k are virtually translated over the calculated displacement vectors $\mathbf{d}(\mathbf{p}_i)$, $\mathbf{d}(\mathbf{p}_j)$, and $\mathbf{d}(\mathbf{p}_k)$, respectively. Using a rasterization algorithm (see *e.g.* Foley *et al.* [22], Chapter 18.7), it is determined which of the pixels \mathbf{p} in the transformed (warped) image $\tilde{I}(\mathbf{p})$ belong to the warped triangle. For every one of those pixels, the inverse displacement $\mathbf{d}^{-1}(\mathbf{p})$ is calculated by linear interpolation of the inverse displacements at the (now displaced) vertices: $\mathbf{d}^{-1}(\mathbf{p}_i + \mathbf{d}(\mathbf{p}_i)) = -\mathbf{d}(\mathbf{p}_i)$. Using this inverse displacement, the grey-value of the warped image \tilde{I} at pixel \mathbf{p} is computed according to the theory described in Section 2:

$$\tilde{I}(\mathbf{p}) = J_r(\mathbf{p})I(\mathbf{p} + \mathbf{d}^{-1}(\mathbf{p})). \quad (12)$$

It can easily be derived that, because of the fact that the inverse displacement of an arbitrary point \mathbf{p} in the image is computed by linear interpolation from the inverse displacements of the three control points \mathbf{p}_i , \mathbf{p}_j , and \mathbf{p}_k of the enclosing triangle Δ_{ijk} , the Jacobian factor J_r in (12) has the form $J_r(\mathbf{p}) = 1 + c(\Delta_{ijk})$, where c is a constant that is dependent only on the enclosing triangle Δ_{ijk} and not on \mathbf{p} itself (see Appendix for details). This implies that J_r is nothing but a constant grey-level scaling factor that needs to be computed only once for every triangle. In order to avoid artifacts at the borders

of triangles, we imposed the additional constraint that the grey-level distribution in the resulting subtraction images must vary in a continuous fashion. Since J_r is a constant within every triangle, this requirement can be satisfied only by taking $c(\Delta_{ijk}) = C$, $\forall \Delta_{ijk}$. In our implementation $C = 0$ (hence, $J_r = 1$), since any other value for C will cause the entire subtraction image to be grey-level scaled, which makes no sense.

In almost all cases, the position $\mathbf{p} + \mathbf{d}^{-1}(\mathbf{p})$ in the original image will not be on the grid. The grey-value at that position can then be obtained by simple bilinear interpolation. In our implementation, the described process of polygon rasterization [22] and texture mapping [26] is carried out real-time by graphics hardware.

Finally, because of the transition from the exposure region to the constant grey-value of the remainder of the image (at the borders), the warping of the corner triangles (as described in Section 4.2) may give rise to very disturbing border artifacts. Since subtraction in those region does not yield any relevant information, the difference values in the corner triangles are explicitly set to zero.

5 Algorithm Overview

In the previous sections, the different operations involved in the registration of two images of a digital angiographic image sequence have been presented and discussed. For clarity, the individual operations are summarized here, together with their parameters.

Given an image sequence of size $M \times M \times N$, the registration of a mask image $I(x, y, 0)$ (without contrasted blood vessels) with respect to the contrast images $I(x, y, n)$, $n \in [1, N - 1] \subset \mathbb{N}$ is accomplished by carrying out the following steps:

1. Calculate the gradient magnitude $\|\nabla L\|$ (at scale σ) of an $M_{\text{edge}} \times M_{\text{edge}}$ sized version of the mask image $I(x, y, 0)$ of the sequence, and extract potential artifact regions by means of thresholding at level Θ_e .

Parameters: M_{edge} , σ , Θ_e .

2. Extract the border $\partial \mathcal{R}_E$ of the exposure region (described by R , X_{\min} , X_{\max} and Y_{\min} , Y_{\max} , see (6)) from the original mask image $I(x, y, 0)$ by analyzing scan-lines from the border to the center of the image, and select a set of border control points.

3. Extract control points from the exposure region \mathcal{R}_E using the thresholded gradient-magnitude $\|\nabla L\|$ version of the mask image $I(x, y, 0)$ and the minimum and maximum distance constraints (D_{\min} , D_{\max} and D_{exp} , based on assumptions about the coherence between neighboring pixels).

Parameters: ϕ_{\min} , ϕ_{\max} , ϕ_{exp} .

4. Given the set of control points $P = \{\mathbf{p}_i\}$ (including four corner points that are positioned outside the image), construct a triangular mesh $\mathcal{D}(P)$ (completely covering the image) using a standard incremental Delaunay triangulation algorithm.

5. For every image $I(x, y, n)$, $n \in [1, N - 1] \subset \mathbb{N}$ in the sequence, calculate the displacement $\mathbf{d}(\mathbf{p}_i, n)$ for the selected control points $\mathbf{p}_i \in P$ (except for the corner points) by maximizing the energy of the histogram of differences, $\mathcal{E}(\mathbf{d})$, in a $W \times W$ neighborhood of these points, using a hill-climbing optimization and using the displacements $\mathbf{d}(\mathbf{p}_i, n - 1)$ of the previous image as an estimate.

Parameter: W .

Sequence	Type	Size
CER	Cerebral	$1024 \times 1024 \times 18$
PER	Peripheral (Femoral)	$512 \times 512 \times 10$
ABD	Abdominal (Kidney)	$1024 \times 1024 \times 6$

Table 1. Specifications of the digital angiographic image sequences used in the experiments as described in Section 6. All data sets were acquired using an Integris V3000 imaging system (Philips Medical Systems, Best, the Netherlands).

Parameter	Value
M_{edge}	512
σ	1.0
Θ_e	15.0
ϕ_{min}	0.04
ϕ_{max}	0.20
ϕ_{exp}	0.01
W	51

Table 2. Values of the parameters of the algorithm during the experiments as described in Section 6. See Section 3 for a detailed description of these parameters.

- For every image $I(x, y, n)$, $n \in [1, N - 1] \subset \mathbb{N}$ in the sequence, warp every triangle Δ_{ijk} in the mesh $\mathcal{D}(P)$ using the displacements $\mathbf{d}(\mathbf{p}_i, n)$, $\mathbf{d}(\mathbf{p}_j, n)$, and $\mathbf{d}(\mathbf{p}_k, n)$ of the constituent control points \mathbf{p}_i and the linearly interpolated displacements at the remaining points $\mathbf{p} \neq \mathbf{p}_i$, using bilinear interpolation of grey-values.

6 Experimental Results

The algorithm as presented in the previous sections was implemented in the C++ programming language [52], utilizing the Open Graphics Library [42]. User-interaction was provided by means of an interface, implemented using Tcl/Tk [44], resulting in a platform-independent application. All experiments were carried out on a relatively low cost O2 workstation (Silicon Graphics, De Meern, the Netherlands) with a 180MHz R5000 IP32 processor, 64MB main memory and 512kB secondary unified instruction/data cache memory, providing special graphics hardware for support of the OpenGL instructions.

The specifications (type and size) of the three data sets (CER, PER, ABD) that were used in order to demonstrate the performance of the algorithm, are presented in Table 1. All data sets are clinical digital angiographic image sequences, acquired on an Integris V3000 imaging system (Philips Medical Systems, Best, the Netherlands). The data sets were manually pre-processed in the sense that uninteresting parts in the images were clipped to a fixed grey-value. During the experiments, the parameters of the algorithm (as summarized in Section 5) were kept fixed to the values in Table 2. For all three image sequences, the first image was taken as the mask image, for which the correspondence was calculated with respect to the subsequent live images.

The results of applying the proposed registration technique to the three data sets of Table 1 are demonstrated in Figures 6, 7, and 8, respectively. In the figures, the original subtraction of the mask image and one of the live images is shown in the top-left panel.

The subtractions after correction for patient motion artifacts, using the proposed approach (based on an irregular triangular mesh), are shown in the bottom-right panel. In order to relate the performance of our algorithm to that of the manual pixel-shifting technique and automatic methods based on regular quadrilateral meshes, the results with these latter methods are shown in the top-right and bottom-left panels, respectively.

In order to give an impression of the speed of the algorithm, the total computation times required to register the image sequences are presented in Table 3. As an example of the distribution of the total computation time over the preprocessing operations and the several images in the sequence, the computation times for data set CER are presented in more detail in Table 4.

7 Discussion

In this section the results obtained by the experiments described in the previous section will be discussed and suggestions for extensions and future research will be presented.

7.1 Discussion of the Results

From Figures 6, 7 and 8 it can be seen that the artifacts in the image sequences cannot be corrected by means of global translation of the mask image, *i.e.*, by applying the pixel-shifting technique provided on standard DSA imaging devices. In all presented sequences, the artifacts were caused by elastic (non-rigid) patient motion, as a result of which registration (using this technique) in one part of the image immediately implied a deterioration of the artifacts in other parts of the image. Furthermore, it can be seen that especially at the borders of the exposure regions, a registration algorithm based on a regular quadrilateral mesh is not capable of completely removing the artifacts. In general, the proposed approach, based on a triangular mesh of irregularly spaced (edge-based) control points yields better registrations and, consequently, better subtractions.

There is a major difference between the registration results of data sets CER and PER (Figures 6 and 7, respectively) and the results of data set ABD (Figure 8) with our method. In the first two sets, the artifacts were removed almost completely, *i.e.*, the algorithm yielded near perfect registrations. In the last data set however, although some artifacts were removed, the registration result still showed some major artifacts. It should be mentioned that these artifacts could not be removed by adjusting one or more of the parameters of the algorithm (Table 2) so as to obtain a larger density of control points. In fact, the artifacts could not even be removed by replacing the hill-climbing optimization by a full-search approach. We will discuss these phenomena in more detail.

The registration result of ABD when using the standard pixel-shifting method (top-right image of Figure 8) reveals that there are parts in the image in which there are several important structures superimposed (*e.g.*, the spine, the catheter and the bowels in the left-middle part of the image). This in contrast with the CER and PER sequences, where the important structures are the result of the projection of only one object in the original three-dimensional scene, *viz.*, bones. In an attempt to remove the artifacts caused by the displacement of the catheter (the black/white curve in the left-middle part of the images of Figure 8), it is inevitable that other artifacts will be introduced (caused by the spine edges in that same region). This phenomenon has been mentioned earlier (in Section 2) as the first limitation of any registration algorithm for projected density images, and explains why even a full search will not be able to retrieve the correspondence.




Figure 6. Registration of data set CER. Top left: the original subtraction of one of the live images from the mask image (showing major motion artifacts). Top right: registration by global translation of the mask image (standard pixel-shifting method) yields only *local* correction (in this case at the bottom right of the image). Bottom left: the registration result using a regular quadrilateral mesh (rubber sheet masking method) still shows some artifacts at the borders of the exposure region. Bottom right: registration using the proposed approach yields a (near) perfect subtraction.

In the center of the image however, the artifacts are caused only by peristaltic motion of the patient. As can be seen from the bottom-right image of Figure 8, our algorithm was able to remove the artifacts near the small vessels (an important region). The larger artifacts on the right side of the image could not be removed. The main reason for this is the increase of noise in the subtraction under translation, causing the match surface to have a local maximum at $\mathbf{d} = (0, 0)$. These artifacts would have been removed if, instead of a hill-climbing optimization, a full search would have been applied. However, this case appeared to be a very exceptional one.




Figure 7. Registration of data set PER. Top left: the original subtraction of one of the live images from the mask image (showing subpixel motion artifacts). Top right: registration by global translation of the mask image (standard pixel-shifting method) yields only *local* correction (in this case at the bottom left of the image). Bottom left: the registration result using a regular quadrilateral mesh (rubber sheet masking method) still shows some minor artifacts at the borders of the exposure region. Bottom right: registration using the proposed approach yields a (near) perfect subtraction.

The computation times presented in Tables 3 and 4 indicate that, in spite of the additional time due to the preprocessing operations (*i.e.*, edge detection, control points selection, triangulation), the proposed approach is faster than commonly used algorithms based on regular quadrilateral meshes. This is mainly due to the edge-based control point selection procedure which, in general, results in a reduction of the number of points for which the displacement needs to be calculated explicitly.

One might argue that algorithms based on regular meshes can be made faster too, by reducing the density of the control points in the mesh. However, this will almost certainly




Figure 8. Registration of data set ABD. Top left: the original subtraction of one of the live images from the mask image (showing major motion artifacts). Top right: registration by global translation of the mask image (standard pixel-shifting method) yields only *local* correction (in this case at the right side of the image). Bottom left: the registration result using a regular quadrilateral mesh (rubber sheet masking method) in which most of the artifacts are still present. Bottom right: registration results using the proposed approach in which there are still major artifacts (see Section 7 for a discussion).

result in a deterioration of the registration accuracy, which is already worse than that of the proposed approach. With the current parameter settings (Table 2), the average number of points in relatively dense regions is about the same with both approaches.

7.2 Extensions and Future Work

After the set of control points P has been tessellated into a Delaunay triangulation $\mathcal{D}(P)$ and the displacement $\mathbf{d}(\mathbf{p}_i)$ has been calculated for every vertex (control point) $\mathbf{p}_i \in P$, it is

Sequence	Algorithm					
	Q-FS	Q-HC	Q-HC-P	T-FS	T-HC	T-HC-P
CER	949.51	31.20	25.96	391.08	19.29	16.42
PER	514.07	12.05	11.91	316.42	8.48	8.40
ABD	341.87	8.77	8.48	131.30	4.66	4.66

Table 3. The total computation times (in seconds) required by the several versions of the algorithm to completely register the data sets presented in Table 1. In this table, Q indicates the use of a regular quadrilateral mesh and T the use of an irregular triangular mesh. The different optimization techniques that were used are full search (FS) (in a range $[-10, 10] \subset \mathbb{N}$ in both x and y direction) and hill climbing (HC). The letter P indicates the application of inter-image displacement prediction.

Image	Algorithm				
	Nr	Q	Q-P	T	T-P
PP		0.00	0.00	0.89	0.91
0		0.00	0.00	0.00	0.00
1		1.65	1.64	0.89	0.92
2		1.81	1.80	1.01	1.00
3		1.85	1.74	1.05	0.90
4		1.83	1.75	1.10	0.95
5		1.90	1.75	1.10	0.96
6		1.86	1.68	1.09	0.90
7		1.87	1.70	1.09	0.87
8		1.86	1.66	1.12	0.88
9		1.83	1.68	1.07	0.90
10		1.81	1.69	1.09	0.93
11		1.83	1.72	1.09	0.93
12		1.84	1.69	1.11	0.92
13		1.88	1.69	1.16	0.93
14		1.84	1.69	1.06	0.86
15		1.83	1.66	1.09	0.89
16		1.85	1.61	1.15	0.90
17		1.86	1.62	1.13	0.87
Total:		31.20	25.96	19.29	16.42

Table 4. The computation times (in seconds) required by the proposed algorithm using an irregular triangular mesh (T) in order to completely register data set CER (see Table 1). The computation times required for a standard rubber-sheet masking approach using a regular quadrilateral mesh (Q) are shown for comparison. PP denotes pre-processing time and the extension P indicates whether inter-image displacement prediction was applied. In all cases, hill-climbing was used for the optimization.

required that the resulting transformation does not cause triangles to fold over. The reason for this requirement is that, in the case of a fold over, the transformation is not consistent in the sense that there exist pixels in the original image that are mapped to the same pixel in the transformed image. Although we did not encounter any occurrences of this phenomenon during the experiments with the data sets in Table 1, a robust implementation

should be able to perform some kind a regularization in such cases. A subject for future research would be the development of a regularization scheme that does not affect the performance of the algorithm in terms of computation times.

Note that the results have been presented on a rather experimental basis. There is no pronounced way in which to quantify the registration results for this type of images. The reason for this is that the actual quality of the resulting subtraction images has to be judged by the person who is responsible for the diagnosis based on these images, *i.e.*, the radiologist. An important activity for the near future will be to carry out a thorough clinical evaluation of the algorithm.

8 Conclusions

In this paper, we have presented a new approach for the registration of digital angiographic image sequences. The method involves the extraction of regions in the image where artifacts can be expected to appear in the case of patient motion. These regions are obtained by thresholding the gradient magnitude of the mask image. Based on assumptions about the coherence of neighboring pixels, a selected set of control points is extracted, for which the displacement is computed explicitly by means of maximizing the energy of the histogram-of-differences similarity measure. A hill-climbing approach is used for the optimization. The complete displacement vector field is constructed from the displacements at these control points, using a Delaunay triangulation and linear interpolation. The final warping of the images is done real-time by graphics hardware.

The overall conclusion from the experimental results is that, in general, the proposed method is effective, very fast, and outperforms algorithms based on regular grids. The best results are obtained in those situations where the important structures in the original three-dimensional scene are exposed in such a way that the grey-level distributions of these structures in the resulting projection images do not overlap. Mostly, this is the case in *e.g.* cerebral and peripheral images. In abdominal images, however, there are often several important structures that can move independently, as a result of which accurate registration with the current approach becomes impossible.

Future research should focus on improving the robustness of the algorithm by including procedures to test for the consistency of the transformation. An important activity for the near future will be to carry out a thorough clinical evaluation of the algorithm.

Appendix

The mapping for which the Jacobian J_r must be computed in (12) is not just the inverse displacement vector field \mathbf{d}^{-1} , but the total reverse mapping \mathbf{d}_r (Theorem 1), defined by

$$\mathbf{d}_r : \mathbb{R}^2 \rightarrow \mathbb{R}^2 : \mathbf{p} \rightarrow \mathbf{p} + \mathbf{d}^{-1}(\mathbf{p}), \quad (13)$$

which we could write as

$$\mathbf{d}_r(x, y) \triangleq \begin{bmatrix} r_x(x, y) \\ r_y(x, y) \end{bmatrix}. \quad (14)$$

The Jacobian of this reverse mapping is then computed as

$$J_r = J_{\mathbf{d}_r} \triangleq \begin{vmatrix} \frac{\partial r_x}{\partial x} & \frac{\partial r_x}{\partial y} \\ \frac{\partial r_y}{\partial x} & \frac{\partial r_y}{\partial y} \end{vmatrix}, \quad (15)$$

where $|\cdot|$ denotes taking the determinant of the matrix.

In general, J_r will have to be computed explicitly for every point in the image (or region of interest). For example, if the region of interest is a rectangle, and the inverse displacement vector field within the region is computed by bilinear interpolation of the inverse displacements of the four corner points (see *e.g.* Mandava *et al.* [39]):

$$\begin{aligned} r_x(x, y) &= x + (a_{10}xy + a_{11}x + a_{12}y + a_{13}), \\ r_y(x, y) &= y + (a_{20}xy + a_{21}x + a_{22}y + a_{23}), \end{aligned} \quad (16)$$

then the Jacobian becomes

$$\begin{aligned} J_r &= \begin{vmatrix} (1 + a_{11} + a_{10}y) & (a_{12} + a_{10}x) \\ (a_{21} + a_{20}y) & (1 + a_{22} + a_{20}x) \end{vmatrix} \\ &= (a_{20} + a_{11}a_{20} - a_{10}a_{21})x + \\ &\quad (a_{10} + a_{22}a_{10} - a_{20}a_{12})y + \\ &\quad 1 + a_{11} + a_{22} + a_{11}a_{22} - a_{12}a_{21}, \end{aligned} \quad (17)$$

which is linearly dependent on x and y . (The eight constant coefficients a_{ij} are easily computed from the eight equations that result when successively substituting the coordinates and displacements of the four corner points into (16).) This result also applies when the image is completely divided into quadrilaterals, which are the suitable polygons in case of a regular grid of control points.

However, in our algorithm the control points are on an *irregular* grid, and they are tessellated into a Delaunay triangulation. The inverse displacement of an arbitrary point in the image is computed from the inverse displacements of the *three* control points constituting the enclosing triangle, by means of *linear* interpolation, *i.e.*

$$\begin{aligned} r_x(x, y) &= x + (a_{11}x + a_{12}y + a_{13}), \\ r_y(x, y) &= y + (a_{21}x + a_{22}y + a_{23}). \end{aligned} \quad (18)$$

It can easily be derived that in this case the Jacobian becomes

$$\begin{aligned} J_r &= \begin{vmatrix} (1 + a_{11}) & a_{12} \\ a_{21} & (1 + a_{22}) \end{vmatrix} \\ &= 1 + a_{11} + a_{22} + a_{11}a_{22} - a_{12}a_{21}, \end{aligned} \quad (19)$$

which is a constant within every triangle. (Again, the coefficients are computed by substituting the coordinates and displacements of the three corner points into (18) and by solving the resulting system of equations.)

At this point, there are two possibilities: either (i) the coefficients a_{ij} are negligible and there is no problem taking $J_r = 1$, or (ii) the coefficients do have significant values and J_r needs to be computed explicitly for every triangle. It is important to note that in the latter case, J_r needs to be computed only once for every triangle, which implies a

great advantage in terms of computational speed as compared to the approach described by *e.g.* Mandava *et al.* [39], where J_r must be computed separately for every point within the region of interest.

It is even more important to have a look at the consequences of including the Jacobian factor into the computations. Since, in our implementation, J_r is a constant within every triangle, J_r is nothing but a constant grey-level scaling factor. Here we have another two possibilities: either (i) the corresponding coefficients a_{ij} in neighboring triangles are almost equal, as a consequence of which there is, again, no problem taking $J_r = 1$, since it would only imply that we do not incorporate a *constant* grey-level scaling of the *entire* image, or (ii) the corresponding coefficients a_{ij} do change substantially from triangle to triangle, and should be recomputed. However, in the latter case, inclusion of the Jacobian factor in the computations will result in substantial grey-level discontinuities at the borders of the triangles, the resulting artifacts of which might even be worse than the ones that we were initially trying to correct for by including this factor. It should be stressed that these artifacts do not only occur in our case of linear interpolation between displacement vectors using (18). They will also occur in the case of bilinear interpolation using (16), which can be seen clearly from the several examples shown by Mandava *et al.* [39].

In order to avoid this type of artifacts, we imposed an additional constraint to the warping algorithm, *viz.*, that the grey-level distribution in the resulting subtraction images should vary in a continuous fashion.⁵ As already pointed out in Section 4.5, this requirement can be satisfied only by explicitly taking J_r to be constant in the entire image. A natural choice is $J_r = 1$ since any other value for J_r will cause the entire subtraction image to be grey-level scaled, which makes no sense.

Acknowledgments

The work described in this paper was carried out at the Image Sciences Institute (ISI), University Medical Center Utrecht (UMCU), the Netherlands, in cooperation with Philips Medical Systems (PMS), Department of X-Ray Diagnostics/Predevelopment, Best, the Netherlands, and was financially supported by the Netherlands Ministry of Economic Affairs. The authors thank Richard Kemkers (PMS) and Wilma Pauw (UMCU) for providing them with the data sets used in the experiments. Onno Wink (ISI) is acknowledged for his contributions in converting the data sets to the local format.

References

- [1] J. K. Aggerwal & N. Nandhakumar, "On the Computation of Motion from Sequences of Images—A Review", *Proceedings of the IEEE*, vol. 76, no. 8, 1988, pp. 917–935.
- [2] R. Althof, M. G. J. Wind, J. T. Dobbins, "A Rapid and Automatic Image Registration Algorithm with Subpixel Accuracy", *IEEE Transactions on Medical Imaging*, vol. 16, no. 1, 1997, pp. 308–316.
- [3] D. I. Barnea & H. F. Silverman, "A Class of Algorithms for Fast Digital Image Registration", *IEEE Transactions on Computers*, vol. 21, no. 2, 1972, pp. 179–186.
- [4] T. Beier & S. Neely, "Feature-Based Image Metamorphosis", *Computer Graphics (SIGGRAPH'92 Conference Proceedings)*, vol. 26, no. 2, 1992, pp. 35–42.
- [5] W. R. Brody, "Hybrid Subtraction for Improved Arteriography", *Radiology*, vol. 141, no. 3, 1981, pp. 828–831.

⁵We realize that it is somewhat sloppy to use the word 'continuous' when dealing with discrete images, but we assume that the reader understands our intentions.

- [6] W. R. Brody, "Digital Subtraction Angiography", *IEEE Transactions on Nuclear Science*, vol. 29, no. 3, 1982, pp. 1176–1180.
- [7] W. R. Brody, D. R. Enzmann, L.-S. Deutsch, A. Hall, N. Pelc, "Intravenous Carotid Arteriography using Line-Scanned Digital Radiography", *Radiology*, vol. 139, no. 2, 1981, pp. 297–300.
- [8] L. G. Brown, "A Survey of Image Registration Techniques", *ACM Computing Surveys*, vol. 24, no. 4, 1992, pp. 325–376.
- [9] T. M. Buzug, C. Lorenz, J. Weese, "Improvement of Vessel Segmentation by Elastically Compensated Patient Motion in Digital Subtraction Angiography Images", in *Computer Analysis of Images and Patterns (CAIP'97)*, G. Sommer, K. Daniilidis, J. Pauli (eds.), vol. 1296 of *Lecture Notes in Computer Science*, Springer-Verlag, Berlin, 1997, pp. 106–113.
- [10] T. M. Buzug & J. Weese, "Improving DSA Images with an Automatic Algorithm based on Template Matching and an Entropy Measure", in *Computer Assisted Radiology (CAR'96)*, H. U. Lemke, M. W. Vannier, K. Inamura, A. G. Farman (eds.), vol. 1124 of *International Congress Series*, Elsevier Science, Amsterdam, 1996, pp. 145–150.
- [11] T. M. Buzug, J. Weese, C. Fassnacht, C. Lorenz, "Image Registration: Convex Weighting Functions for Histogram-Based Similarity Measures", in *CVRMed-MRCAS'97*, J. Troccaz, E. Grimson, R. Mösges (eds.), vol. 1205 of *Lecture Notes in Computer Science*, Springer-Verlag, Berlin, 1997, pp. 203–212.
- [12] J. F. Canny, "A Computational Approach to Edge Detection", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, no. 6, 1986, pp. 679–698.
- [13] J. Y. Chiang & B. J. Sullivan, "Coincident Bit Counting—A New Criterion for Image Registration", *IEEE Transactions on Medical Imaging*, vol. 12, no. 1, 1993, pp. 30–38.
- [14] W. A. Chilcote, M. T. Modic, W. A. Pavlicek, J. R. Little, A. J. Furian, P. M. Duchesneau, M. A. Weinstein, "Digital Subtraction Angiography of the Carotid Arteries: A Comparative Study in 100 Patients", *Radiology*, vol. 139, no. 2, 1981, pp. 287–295.
- [15] G. S. Cox, "Review: Template Matching and Measures of Match in Image Processing", Tech. Rep., University of Cape Town, Cape Town, 1995.
- [16] G. S. Cox & G. de Jager, "Automatic Registration of Temporal Image Pairs for Digital Subtraction Angiography", in *Medical Imaging: Image Processing*, M. H. Loew (ed.), vol. 2167 of *Proceedings of SPIE*, The International Society for Optical Engineering, Bellingham, WA, 1994, pp. 188–199.
- [17] L. S. Davis, Z. Wu, H. Sun, "Contour-Based Motion Estimation", *Computer Vision, Graphics and Image Processing*, vol. 23, no. 3, 1983, pp. 313–326.
- [18] J. M. Fitzpatrick, "The Existence of Geometrical Density-Image Transformations Corresponding to Object Motion", *Computer Vision, Graphics and Image Processing*, vol. 44, no. 2, 1988, pp. 155–174.
- [19] J. M. Fitzpatrick, J. J. Grefenstette, D. R. Pickens, M. Mazer, J. M. Perry, "A System for Image Registration in Digital Subtraction Angiography", in *Image Processing in Medical Imaging*, C. N. de Graaf & M. A. Viergever (eds.), Plenum Press, New York, NY, 1988, pp. 415–435.
- [20] J. Flusser, "An Adaptive Method for Image Registration", *Pattern Recognition*, vol. 25, no. 1, 1992, pp. 45–54.
- [21] S. V. Fogel, "The Estimation of Velocity Vector Fields from Time-Varying Image Sequences", *CVGIP: Image Understanding*, vol. 53, no. 3, 1991, pp. 253–287.
- [22] J. Foley, A. van Dam, S. K. Feiner, J. F. Hughes, *Computer Graphics: Principles and Practice*, 2nd ed., Systems Programming Series, Addison-Wesley, Reading, MA, 1990.
- [23] A. Goshtasby, "Piecewise Linear Mapping Functions for Image Registration", *Pattern Recognition*, vol. 19, no. 6, 1986, pp. 459–466.
- [24] A. Goshtasby, "Piecewise Cubic Mapping Functions for Image Registration", *Pattern Recognition*, vol. 20, no. 5, 1987, pp. 525–533.
- [25] A. Goshtasby, G. C. Stockman, C. V. Page, "A Region Based Approach to Digital Image Registration with Subpixel Accuracy", *IEEE Transactions on Geoscience and Remote Sensing*, vol. 24, no. 3, 1986, pp. 390–399.
- [26] P. S. Heckbert, "Survey of Texture Mapping", *IEEE Computer Graphics and Applications*, vol. 6, no. 11, 1986, pp. 56–67.
- [27] E. C. Hildreth, "The Detection of Intensity Changes by Computer and Biological Vision Systems", *Computer Vision, Graphics and Image Processing*, vol. 22, no. 1, 1983, pp. 1–27.

- [28] E. C. Hildreth, "The Computation of the Velocity Field", *Proceedings of the Royal Society of London, Series B*, vol. 221, 1984, pp. 189–220.
- [29] B. J. Hillman, T. W. Ovitt, S. Nudelman, H. D. Fischer, M. M. Frost, P. Capp, H. Roehrig, G. Seeley, "Digital Video Subtraction Angiography of Renal Vascular Abnormalities", *Radiology*, vol. 139, no. 2, 1981, pp. 277–280.
- [30] B. K. P. Horn & B. G. Schunck, "Determining Optical Flow", *Artificial Intelligence*, vol. 17, 1981, pp. 185–203.
- [31] P. Hua & I. Fram, "Feature-Based Image Registration for Digital Subtraction Angiography", in *Image Processing*, M. H. Loew (ed.), vol. 1898 of *Proceedings of SPIE*, The International Society for Optical Engineering, Bellingham, WA, 1993, pp. 24–31.
- [32] R. A. Kruger, F. J. Miller, J. A. Nelson, P. Y. Liu, W. Bateman, "Digital Subtraction Angiography using a Temporal Bandpass Filter: Associated Patient Motion Properties", *Radiology*, vol. 145, no. 2, 1982, pp. 315–320.
- [33] R. A. Kruger, C. A. Mistretta, A. B. Crummy, J. F. Sackett, M. M. Goodsitt, S. J. Riederer, T. L. Houk, C.-G. Shaw, D. Flemming, "Digital K -Edge Subtraction Radiography", *Radiology*, vol. 125, no. 1, 1977, pp. 243–245.
- [34] R. A. Kruger, J. A. Nelson, D. G. Roy, F. J. Miller, R. E. Anderson, P. Liu, "Dynamic Tomographic Digital Subtraction Angiography using Temporal Filtration", *Radiology*, vol. 147, no. 3, 1983, pp. 863–867.
- [35] R. A. Kruger, M. Sedaghati, D. G. Roy, P. Liu, J. A. Nelson, W. Kubal, P. Del Rio, "Tomosynthesis Applied to Digital Subtraction Angiography", *Radiology*, vol. 152, no. 3, 1984, pp. 805–808.
- [36] C. L. Lawson, "Software for C^1 Surface Interpolation", in *Mathematical Software III*, J. R. Rice (ed.), Academic Press, New York, NY, 1977, pp. 161–194.
- [37] D. T. Lee & B. J. Schachter, "Two Algorithms for Constructing a Delaunay Triangulation", *International Journal of Computer and Information Sciences*, vol. 9, no. 3, 1980, pp. 219–242.
- [38] J. B. A. Maintz & M. A. Viergever, "A Survey of Medical Image Registration", *Medical Image Analysis*, vol. 2, no. 1, 1998, pp. 1–36.
- [39] V. R. Mandava, J. M. Fitzpatrick, D. R. Pickens, "Adaptive Search Space Scaling in Digital Image Registration", *IEEE Transactions on Medical Imaging*, vol. 8, no. 3, 1989, pp. 251–262.
- [40] D. Marr & E. C. Hildreth, "Theory of Edge Detection", *Proceedings of the Royal Society of London, Series B*, vol. 207, 1980, pp. 187–217.
- [41] C. A. Mistretta, M. G. Ort, F. Kelcz, J. R. Cameron, M. P. Sieband, A. B. Crummy, "Absorption Edge Fluoroscopy using Quasimonoeenergetic X-Ray Beams", *Investigative Radiology*, vol. 8, no. 6, 1973, pp. 402–412.
- [42] J. Neider, T. Davis, M. Woo, *OpenGL Programming Guide*, Addison-Wesley, Reading, MA, 1995.
- [43] H. Oung & A. M. Smith, "Real Time Motion Detection in Digital Subtraction Angiography", in *Proceedings of the International Symposium on Medical Images and Icons*, A. Deurinckx, M. H. Loew, J. M. S Prewitt (eds.), IEEE Computer Society Press, Silver Spring, RI, 1984, pp. 336–339.
- [44] J. K. Ousterhout, *Tcl and the Tk Toolkit*, Professional Computing Series, Addison-Wesley, Reading, MA, 1994.
- [45] M. J. Potel & D. E. Gustafson, "Motion Correction for Digital Subtraction Angiography", in *Frontiers of Engineering and Computing in Health Care: Proceedings of the 5th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, G. C. Gerhard & W. T. Miller (eds.), 1983, pp. 166–169.
- [46] M. J. D. Powell, "An Efficient Method for Finding the Minimum of a Function of Several Variables without Calculating Derivatives", *The Computer Journal*, vol. 7, 1964, pp. 155–162.
- [47] W. K. Pratt, "Correlation Techniques of Image Registration", *IEEE transactions on Aerospace and Electronic Systems*, vol. 10, 1974, pp. 353–358.
- [48] D. Ruprecht, *Geometrische Deformationen als Werkzeug in der Graphischen Datenverarbeitung*, Ph.D. thesis, Universität Dortmund, Dortmund, 1994.
- [49] D. Ruprecht & H. Müller, "Image Warping with Scattered Data Interpolation", *IEEE Computer Graphics and Applications*, vol. 15, no. 2, 1995, pp. 37–43.

- [50] J. Shi & C. Tomasi, “Good Features to Track”, in *IEEE Conference on Computer Vision and Pattern Recognition*, 1994, pp. 593–600.
- [51] G. C. Stockman, S. Kopstein, S. Benett, “Matching Images to Models for Registration and Object Detection via Clustering”, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 4, no. 3, 1982, pp. 229–241.
- [52] B. Stroustrup, *The C++ Programming Language*, 2nd ed., Addison-Wesley, Reading, MA, 1991.
- [53] R. Szeliski & J. Coughlan, “Spline-Based Image Registration”, *International Journal of Computer Vision*, vol. 22, no. 3, 1997, pp. 199–218.
- [54] C. Tomasi & T. Kanade, “Shape and Motion from Image Streams: A Factorization Method—Part 3. Detection and Tracking of Point Features”, Tech. Rep. CMU-CS-91-132, Carnegie Mellon University, Pittsburgh, 1991.
- [55] P. A. van den Elsen, E.-J. D. Pol, M. A. Viergever, “Medical Image Matching—A Review with Classification”, *IEEE Engineering in Medicine and Biology*, vol. 12, no. 1, 1993, pp. 26–39.
- [56] L. van Tran & J. Sklansky, “Flexible Mask Subtraction for Digital Angiography”, *IEEE Transactions on Medical Imaging*, vol. 11, no. 3, 1992, pp. 407–415.
- [57] A. Venot, J. F. Lebruchec, J. C. Roucayrol, “A New Class of Similarity Measures for Robust Image Registration”, *Computer Vision, Graphics and Image Processing*, vol. 28, no. 2, 1984, pp. 176–184.
- [58] A. Venot & V. Leclerc, “Automated Correction of Patient Motion and Gray Values Prior to Subtraction in Digitized Angiography”, *IEEE Transactions on Medical Imaging*, vol. 3, no. 4, 1984, pp. 179–186.
- [59] A. Venot, L. Pronzato, E. Walter, “Comments about the Coincident Bit Counting (CBC) Criterion for Image Registration”, *IEEE Transactions on Medical Imaging*, vol. 13, no. 3, 1994, pp. 565–566.
- [60] L. A. J. Verhoeven, *Digital Subtraction Angiography. The Technique and an Analysis of the Physical Factors Influencing the Image Quality*, Ph.D. thesis, Delft University of Technology, Delft, 1985.
- [61] D. F. Watson, “Computing the n -Dimensional Delaunay Tessellation with Application to Voronoi Polytopes”, *The Computer Journal*, vol. 24, no. 2, 1981, pp. 167–172.
- [62] D. F. Watson & G. M. Philip, “Survey: Systematic Triangulations”, *Computer Vision, Graphics and Image Processing*, vol. 26, no. 2, 1984, pp. 217–223.
- [63] G. Wolberg, *Digital Image Warping*, IEEE Computer Society Press, WA, 1990.
- [64] M. Yanagisawa, S. Shigemitsu, T. Akatsuka, “Registration of Locally Distorted Images by Multi-window Pattern Matching and Displacement Interpolation: The Proposal of an Algorithm and Its Application to Digital Subtraction Angiography”, in *Proceedings of the Seventh International Conference on Pattern Recognition*, M. D. Levine (ed.), vol. 2, IEEE Publishing Services, New York, NY, 1984, pp. 1288–1291.
- [65] K. J. Zuiderveld, B. M. ter Haar Romeny, M. A. Viergever, “Fast Rubber Sheet Masking for Digital Subtraction Angiography”, in *Science and Engineering of Medical Imaging*, M. A. Viergever (ed.), vol. 1137 of *Proceedings of SPIE*, The International Society for Optical Engineering, Bellingham, WA, 1989, pp. 22–30.