OPTICAL FLOW TO ANALYSE STABILISED IMAGES OF THE BEATING HEART

Martin Gröger
Institute of Robotics and Mechatronics, German Aerospace Center (DLR)
82234 Oberpfaffenhofen

Gerd Hirzinger
Institute of Robotics and Mechatronics, German Aerospace Center (DLR)
82234 Oberpfaffenhofen

Keywords: Optical flow, motion field, image stabilisation, motion compensation, tracking, beating heart surgery, minimally invasive surgery, image-guided surgery, medical robotics.

Abstract: An optical flow method is developed to analyse the motion of the beating heart surface and the efficacy of strategies to stabilise this motion. Although reduced by mechanical stabilisers, residual tissue motion makes safe surgery still difficult and time consuming. Compensation for this movement is therefore highly desirable. Images of the heart surface, captured by a video endoscope, can be further stabilised based on motion information obtained by tracking natural landmarks in realtime. The remaining motion on the heart surface is assessed by a specially developed optical flow approach: It estimates the image velocities based on a robust region-based strategy and provides a reliable measure of the motion field of the heart. The analysis shows that tissue motion can be reduced by a global motion correction strategy while local motion differences remain.

1 INTRODUCTION

Intraoperative organ motion, induced by heart beat and respiration poses special requirements to robot-assisted surgery. Although the motion of the beating heart is mechanically stabilised, significant residual motion remains, which makes safe surgery difficult and time consuming. Therefore, motion compensation is a highly desired issue, particularly in minimally invasive beating heart surgery (Jacobs et al., 2003). Recognition of this motion, captured by a video endoscope, is essential for a motion compensation system, in which the images of the beating heart appear stabilised to the surgeon and his instruments are moved accordingly by the robot.

A motion correction strategy is introduced to reduce the heart motion remaining in the image. It is based on realtime tracking of natural landmarks on the heart surface (Gröger et al., 2002). The effectiveness of motion correction is analysed by a specially designed region-based optical flow method. It is also based on the mentioned tracking approach.

Estimating the motion of the heart by natural landmarks instead of artificial markers (Nakamura et al., 2001; Ginhoux et al., 2004; Gröger et al., 2004) is especially attractive since no markers need to be introduced into the field of surgery. Moreover, region-based tracking of natural landmarks yields a particular texture unique for each landmark. This easily allows to track several landmarks concurrently, whereas using identical artifical landmarks bears the danger of ambiguities.

Optical flow is the apparent motion of brightness patterns in an image (Horn, 1986). The motion field is a purely geometric concept, which assigns a velocity vector to each point in an image. The 2D motion field is the projection of the 3D velocities of surface points onto the image plane. The goal of optical flow estimation is to compute an approximation to the 2D motion field from spatiotemporal patterns of image intensity. Ideally, the optical flow corresponds to the motion field (Horn, 1986).

Different approaches exist to compute the optical flow (Barron et al., 1994), among them differential and region-based methods. Differential techniques, as proposed by (Horn and Schunck, 1981) and (Lucas and Kanade, 1981), compute velocity from spatiotemporal derivatives of image intensity (Barron et al., 1994). These techniques calculate the optical flow based on the gradient constraint equation, which is defined as follows for a two-dimensional image $I \stackrel{\text{def}}{=} I(x, y, t)$ with coordinates (x, y) and time t

$$I_x u + I_y v + I_t = 0, \tag{1}$$

where I_x, I_y, and I_t denote the partial derivatives of

Gröger M. and Hirzinger G. (2006).
OPTICAL FLOW TO ANALYSE STABILISED IMAGES OF THE BEATING HEART.
In *Proceedings of the First International Conference on Computer Vision Theory and Applications*, pages 237-244
DOI: 10.5220/0001378602370244

I with respect to x, y, and t. Since this linear equation has two unknowns, i.e. the components u and v of the motion vector at a particular image position in x and y directions, further constraints are necessary (Barron et al., 1994). The components u and v of image velocity can be constrained by combining local estimates of velocity through space and time. The method by (Horn and Schunck, 1981) uses global smoothness constraints (regularisation) in which the motion field is defined implicitly in terms of the minimum of a functional defined over the image. The method by (Lucas and Kanade, 1981) fits the measurements in each neighbourhood to a local model for the 2D velocity, using least-squares minimisation.

However, differential methods pose special requirements, which can limit the detected image velocities to 1 px/frame. Moreover, the smoothness constraint incorporates assumptions on the homogeneity of the motion field, which can change the estimated value of a particular image velocity undesiredly.

Region-based optical flow techniques do not require global assumptions on the imaging data and also deal well with large image velocities, which is required for the analysis of the given sequences of the beating heart. Therefore, a region-based optical flow strategy is designed.

The approach of motion tracking by natural landmarks on the heart surface is introduced first. Based on it, a method for optical flow computation is developed, together with a measure of the resulting motion fields. Furthermore, a global motion correction scheme, also based on the motion tracking approach, is introduced. It is evaluated by experiments applying the proposed optical flow method to synthetic, real and motion corrected image sequences of the heart.

2 METHODS

2.1 Motion Tracking

Motion of the beating heart is captured by tracking natural landmarks on the heart surface recorded by a video endoscope (Fig. 1).

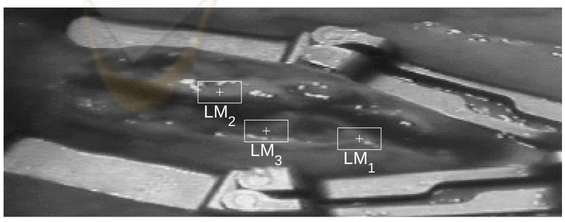

Figure 1: Heart with mechanical stabiliser and landmarks.

However, specular reflections of the light source on the glossy surface of the beating heart (Fig. 1) can disturb tracking of natural landmarks considerably, which makes their appropriate treatment a prerequisite. Tracking outliers caused by specular reflections on the heart surface are detected and eliminated by the method presented in (Gröger et al., 2005), which substitutes the reflections by intensity information taking into account local image structure. This scheme is applied to images of the beating heart prior to tracking.

Natural landmarks are used to track local motion on the beating heart as described in (Gröger et al., 2002). These landmarks are selected image patterns of a given size and offer a particular texture. Tracking is performed by intensity based pattern matching of a given reference landmark in a particular search area. The sum of squared differences (SSD) error measure is used to find the position of the best matching pattern. The nonlinear distortions of the heart surface can be approximated in the 2D image plane by an affine motion model, with the two translational parameters being dominant such that they are sufficient for tracking in the mechanically stabilised area on the beating heart (Gröger et al., 2002).

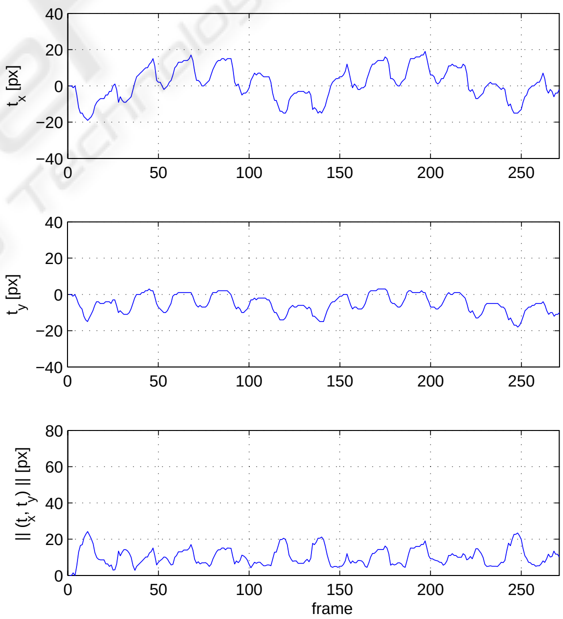

Figure 2: Motion trajectory of landmark LM_1 (translation in x and y directions and overall ($\| (t_x, t_y) \| \stackrel{\mathrm{def}}{=} \sqrt{t_x^2 + t_y^2}$)).

Motion tracking for the analyses below is performed in the translational search space, searched exhaustively for the best matching pattern.

Three example landmarks are shown in Fig. 1. A block size of 30 px $\times$ 15 px proves appropriate for the heart sequences investigated below, while the translational search area can be restricted to 70 px $\times$ 50 px.

Other tracking approaches update the reference pattern continually, e.g. from frame to frame. This bears the danger of tracking drift (current reference becomes more and more dissimilar from the original reference) and total loss of the pattern.

The proposed tracking strategy keeps the same reference pattern throughout the image sequence. This makes tracking particularly robust, since the loss of the pattern in one frame does not imply the total loss of the pattern for the remaining sequence. As a drawback, this strategy increases the search space to find subsequent patterns. However, since the motion of the heart, induced by heart beat and respiration, is quasi-periodic, the range of motion is restricted (see motion trajectories in Fig. 2).

2.2 Optical Flow Calculation

A suitable optical flow strategy is required to analyse the motion of the beating heart and to assess the effectiveness of the motion correction strategy. While reliability is most important for motion analysis, the strategy does not need to run in realtime.

Magnitudes of motion of the beating heart can be rather large, especially when tracking by referring to the original reference image as described above, leading to magnitudes of up to 40 px for the given image sequence (Fig. 2). As the proposed motion correction also keeps a reference image, the optical flow method designed for its evaluation should calculate the velocities between the current and the reference image.

Since differential motion estimation techniques are often restricted to small motion displacements (Barron et al., 1994), e.g. less than 1 px, satisfying results of motion estimation cannot be expected in this context. Results for the optical flow techniques by (Horn and Schunck, 1981) and (Lucas and Kanade, 1981) are included in the analysis below.

Therefore, a region-based optical flow method is proposed to analyse the motion of the beating heart. The method uses the motion tracking strategy described above applied to every image pixel to build a dense motion field. Outliers by specular reflections on the heart surface are avoided by applying the described elimination strategy prior to tracking.

2.3 Optical Flow Measure

The optical flow algorithm yields a dense motion field of the beating heart surface. It can be represented by a needle map, such as in Fig. 9, which, however, does not show the motion vector of every single pixel, for reasons of clarity. The question arises how these motion fields can be assessed to allow for evaluation and comparison of particular heart motion scenarios, e.g. after application of different motion correction algorithms. In the following, an appropriate measure of optical flow is presented, together with a strategy to ensure its quality by outlier removal.

The magnitude of the remaining motion is especially important for assessing the performance of motion correction. The direction of the remaining motion is not so critical, since the goal of motion correction is to minimise the remaining motion such that its magnitude becomes as small as possible. Moreover, direction is not well defined for good motion correction with remaining speeds approaching zero (also see the corresponding frequency distribution in Fig. 4 below, which lacks compactness). Poor motion correction, however, is characterised by large magnitudes of remaining motion, while the directions do not provide additional information. Therefore, only the magnitude of motion is considered for the evaluation of heart motion. The mean magnitude of all vectors of a motion field M is proposed as a global measure of motion and denoted as the *mean speed* of M

$$\mu\,(M) \stackrel{\text{def}}{=} \frac{1}{|M|} \sum_{m \in M} \parallel m \parallel . \qquad (2)$$

Outliers in the motion field, however, can affect the quality of this measure. Therefore, they should be detected and excluded from the motion field. The motion fields in Fig. 11 show areas of outliers. Possible outliers caused by specular reflections are prevented by the proposed motion estimation strategy which includes their elimination. Since the motion analyses are focussed on the beating heart surface, the image areas of the mechanical stabiliser are excluded, the glossy surface of which often does not bear sufficient texture for tracking. Further, there can still be areas on the heart surface without significant texture, i.e. rather homogeneous regions, in which robust tracking is not possible and outliers occur.

Motion fields of the beating heart show a clear homogeneity of velocity, which can be analysed by the corresponding histograms. Figures 3 and 4 show frequency distributions of speed and direction for original and motion corrected image pairs. The concentration around the maximum frequency in the speed histograms confirms the observed homogeneity of the motion field. Outliers are characterised by disturbing this homogeneity and appear as less frequent values farther from the maximum. Therefore, outliers can be detected by thresholds in the frequency distribution.

Constant thresholds for a longer image sequence of the beating heart, however, are not recommended, since the frequency distribution of these images is too broadly spread to find such thresholds (Fig. 5). This is due to the periodically changing motion direction of the heart in the image plane and the corresponding speed fluctuations. Therefore, individual thresholds for outlier detection are calculated for each image.

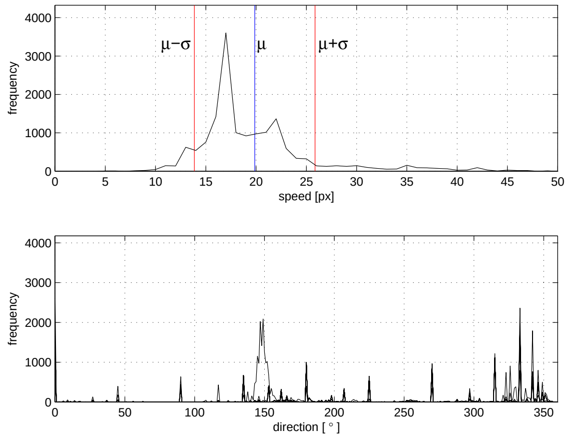

Figure 3: Frequency distributions of speed and direction in an uncorrected image pair (frame 14).

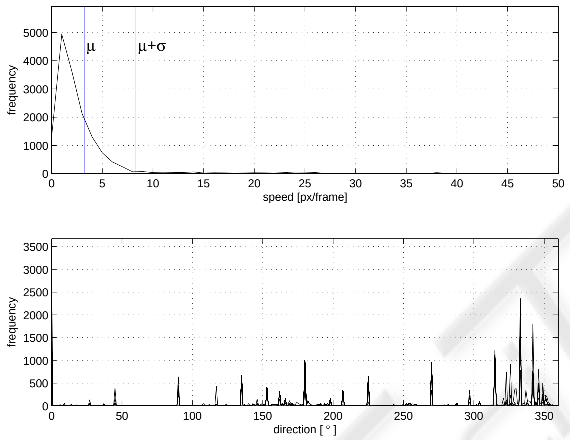

Figure 4: Frequency distributions of speed and direction in a motion corrected image (frame 14, global motion correction by landmark LM_2).

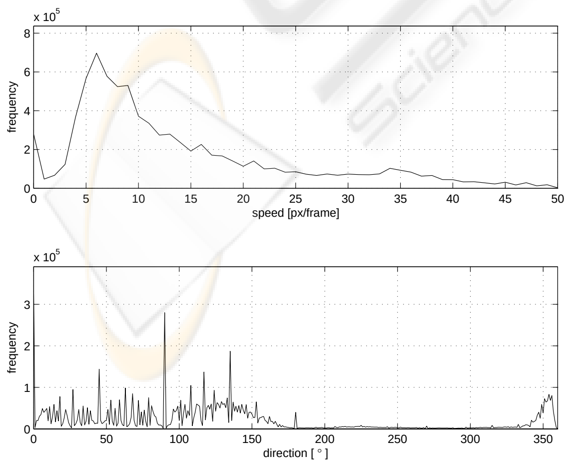

Figure 5: Frequency distributions of speed and direction in 271 images of an uncorrected image sequence.

Making the thresholds dependent on the mean speed μ of the current image yields good outlier detection. Let σ be the standard deviation of speeds of the motion field. As shown in Figures 3 and 4 an area of acceptance of $\mu \pm \sigma$ captures the most frequent and nearby speeds well, particularly for lower speeds as occurring in motion compensated image sequences. Figures 6 and 7 show the removal of outliers for examples in the original and motion corrected image sequences. The $\mu \pm \sigma$ threshold rejects 5.5 % of all motion vectors as outliers in the motion corrected image sequence but 18.5 % in the original one. This fact and the corresponding histogram distribution indicate that a $\mu \pm 2\sigma$ threshold could be more appropriate for the original image sequence with higher speeds, which only rejects 6.0 % of all motion vectors as outliers. This threshold reduces the number of outliers to 4.0 % in motion corrected images. However, since a large area of acceptance bears the danger of undetected outliers and the focus of analysis is on motion corrected images, the stricter threshold of $\mu \pm \sigma$ is used in the following to measure the mean speed of motion fields of the beating heart. So a motion vector m is accepted only if its speed is in the $\mu \pm \sigma$ area, i.e.

$$\parallel m \parallel \in [\, \mu - \sigma \, , \, \mu + \sigma \,] \, . \qquad (3)$$

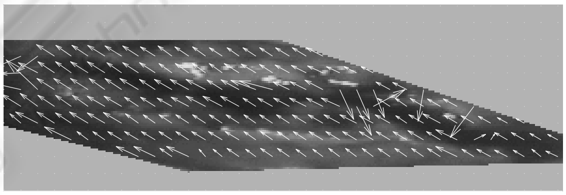

(a) Without removal of outliers.

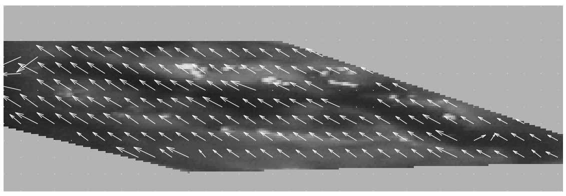

(b) Acceptance region $\mu \pm 2\sigma$, 6.0 % outliers.

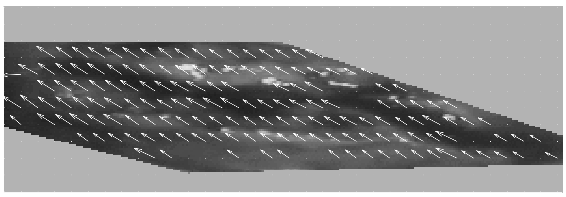

(c) Acceptance region $\mu \pm \sigma$, 18.5 % outliers.

Figure 6: Motion fields with outliers removed by thresholding (frame 14 in original image sequence).

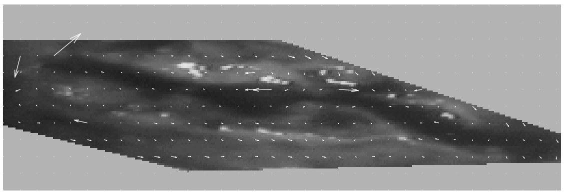

(a) Without removal of outliers.

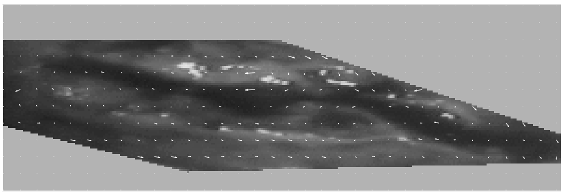

(b) Acceptance region $\mu \pm 2\sigma$, 4.0 % outliers.

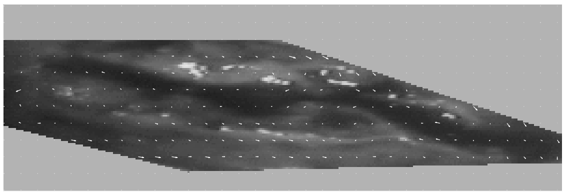

(c) Acceptance region $\mu \pm \sigma$, 5.5 % outliers.

Figure 7: Motion fields with outliers removed by thresholding (frame 14 in motion corrected image sequence).

2.4 Motion Correction

Stabilisation of images of the beating heart uses the tracked motion of natural landmarks. A global motion vector is used to correct the motion by translation constant all over the image. A first strategy uses the motion of a single landmark to build the motion correction vector. Combining the motion of several landmarks increases robustness. A given tracked landmark i with reference position $(x_{i,0}, y_{i,0})$ yields a 2D motion vector c_i in the image plane at time t, i.e.

$$c_i \stackrel{\text{def}}{=} \begin{pmatrix} u_i \\ v_i \end{pmatrix} = \begin{pmatrix} x_{i,t} - x_{i,0} \\ y_{i,t} - y_{i,0} \end{pmatrix}. \tag{4}$$

Motion is corrected by translating the whole image by the vector $-c_i$. This strategy can easily be extended to several landmarks. The global motion vector c is calculated by the mean of n given landmarks with motion vectors c_i, i.e.

$$c \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^{n} c_i. \tag{5}$$

Using the mean motion vector of several landmarks makes the resulting global motion correction vector c more characteristic of the motion of the heart surface and increases robustness in case of tracking outliers.

This motion correction strategy is similar the visual servoing approach of tracking a particular landmark on the heart surface by the viewing camera, which only allows for global motion correction in the image.

3 EXPERIMENTS

The methods of optical flow and the mean speed measure are used to analyse motion in image sequences of the beating heart and to assess the developed strategy for image stabilisation by global motion correction.

First, optical flow is verified on a synthetic image sequence of the heart, for which the underlying motion is known. Then the optical flow strategy is used to estimate the motion of real image sequences of the beating heart – first without motion correction, then with the global motion correction strategy applied.

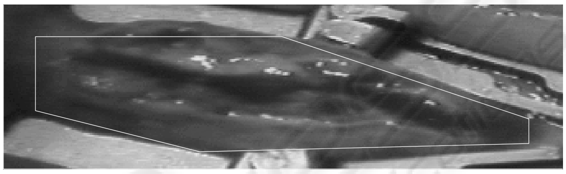

Figure 8: Region of interest (ROI) on the beating heart.

Since motion analysis focuses on the heart surface, the optical flow measure is only calculated in the region of interest (ROI) between the branches of the mechanical stabiliser as shown in Fig. 8.

3.1 Synthetic Image Sequence

For an objective, quantitative analysis of optical flow algorithms the true motion is desired to be known in advance. Since, however, this is not the case for the heart sequences, a synthetic sequence with known motion is created. Of course, this sequence should have properties close to the real image sequence.

Figure 9: Motion field in the region of interest of the synthetical image sequence with constant global motion.

The prevailing motion in image sequences of the mechanically stabilised beating heart is translational (Gröger et al., 2002). Moreover, the analysis of trajectories of several natural landmarks on the beating heart, e.g. those in Fig. 2, shows that the occurring motion can be rather large, up to 40 px. Therefore, a synthetic image sequence is created by translating one reference image of the heart by up to 20 px in steps of 0.2 px using bilinear interpolation. The maximum motion of 20 px occurs when comparing the last frame to the first one. Figure 9 shows the motion field of this sequence.

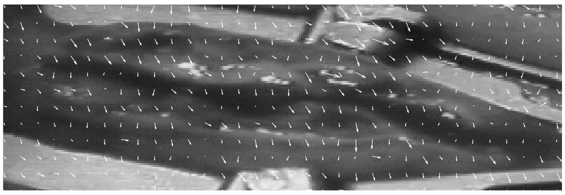

(a) Horn-Schunck ($\lambda = 0.1$)

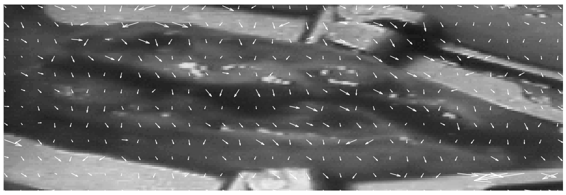

(b) Lucas-Kanade (window 15 px $\times$ 15 px)

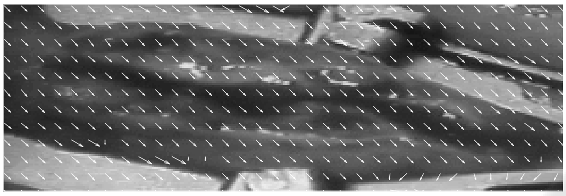

(c) Block-Matching (block size 30×15, precision 0.125 px)

Figure 10: Calculated motion fields of synthetic image sequence (image translated by 1.6 px, vectors scaled by 2).

Two differential optical flow strategies, (Horn and Schunck, 1981) and (Lucas and Kanade, 1981), and the proposed region-based strategy are investigated on this sequence. Figure 10 shows the result for the image pair being translated by 1.6 px. The first differential strategy (Horn and Schunck, 1981) uses a smoothness parameter of $\lambda = 0.1$, while the second approach (Lucas and Kanade, 1981) controls smoothness locally by a window of size 15 px $\times$ 15 px. The block-matching strategy uses a 30 px $\times$ 15 px block and is run at subpixel precision of 0.125 px.

As can be seen in Fig. 10, only the block-matching strategy is able to detect the true motion of the image sequence reliably for a translation of 1.6 px. The results at higher translations are even much worse for the two differential approaches, which generally have problems to detect larger speeds (Barron et al., 1994).

Therefore, the block-matching approach is used in the following to calculate the optical flow measures of the beating heart.

3.2 Real Image Sequence

The proposed region-based optical flow approach is applied to a real image sequence of the beating heart, consisting of 271 frames at a framerate of 25 Hz.

Since the true motion of the real image sequence is unknown, the optical flow algorithm has to be evaluated by special measures on the motion field itself, without the possibility of comparing the result to the underlying motion. The issues of homogeneity and outliers are particularly important in this context.

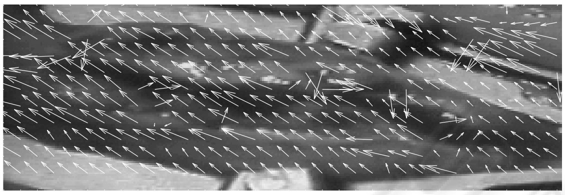

(a) Without elimination of specular reflections.

(b) With elimination of specular reflections.

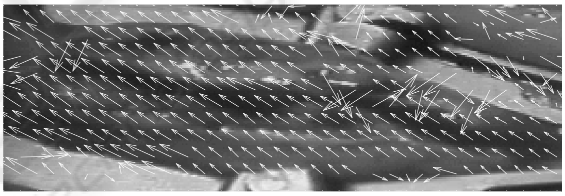

(c) With elimination of specular reflections and illumination compensation.

Figure 11: Motion fields of the real image sequence calculated by block-matching approach (frame 13, block size 30 px $\times$ 15 px, vector lengths scaled by 0.5).

Figure 11 shows the motion fields of the original image sequence calculated by different versions of the region-based optical flow approach for a selected image pair. The application of specular reflection removal prior to the block-matching developed in (Gröger et al., 2005) is necessary to avoid outliers. Also, the illumination compensation scheme, providing mean compensated patterns when finding the best match as described in (Gröger et al., 2002) is important. The remaining outliers in the motion field occur on the one hand at the mechanical stabiliser, which is outside the ROI (see Fig. 8), and on the other hand e.g. at image areas with insufficient texture information and are therefore not considered in the proposed optical flow measure of mean speeds (see Sect. 2.3).

The motion fields calculated by the region-based optical flow strategy are homogeneous and correspond to the impression of motion with visual inspection. This confirms the result obtained by the synthetic image sequence in the previous section that this approach is well suited to detect the motion of the beating heart in the given image sequences.

The mean speed of all motion vectors in the region of interest and of the whole image sequence is 10.21 px/frame; its mean speed per image over this sequence can be seen in Fig. 12.

3.3 Motion Corrected Sequence

The described global motion correction scheme is applied to the given image sequence of the beating heart. The results of motion correction are evaluated by the mean speed of the remaining motion field.

Trajectories of different landmarks are used to perform motion correction. First, the result of motion correction by single landmarks is shown. Then motion information from several landmarks is combined, and finally also landmarks on the mechanical stabiliser are used. The evaluation of motion always relates to the region of interest (ROI) inside the branches of the mechanical stabiliser (see Fig. 8).

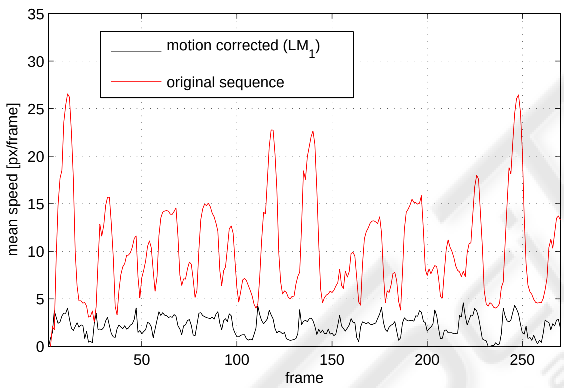

Figure 12: Mean speeds of original and corrected image sequence (global translation by motion of landmark LM_1).

Motion correction based on single landmarks is investigated for the three landmarks given in Fig. 1. Figure 12 provides the mean speed of motion in the ROI in the original image sequence and after image correction by the motion of landmark LM_1. It shows that image motion is significantly reduced and also stays below a certain threshold of 5 px/frame. The mean speed of the whole sequence is 10.21 px/frame for the original image sequence, but can be reduced to speeds between 2 and 3 px/frame by motion correction depending on the selected landmark (see Tab. 1). This also shows that the quality of motion correction depends on the selected landmark. To make motion correction less dependent on single landmarks and thus make the calculated global motion correction vector more characteristic of the occurring heart motion, the mean motion of several landmarks is used as introduced in Sect. 2.4.

Table 1: Mean speeds of image sequence, motion globally corrected by single landmarks.

landmark	mean speed
LM_1	2.15 px/frame
LM_2	2.72 px/frame
LM_3	2.25 px/frame

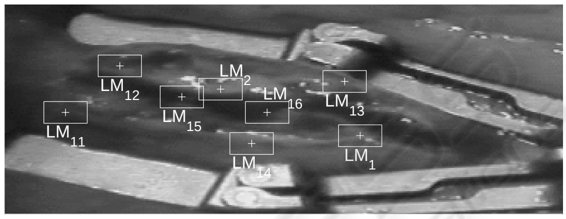

Figure 13: Heart surface with landmark choice C_5.

Image correction by the mean motion of several landmarks is investigated for the eight landmarks shown in Fig. 13. An increasing number of landmarks is used to correct the motion of the beating heart. The results in Tab. 2 show that the mean speed can be reduced from 2.15 px/frame for a single landmark down to 1.91 px/frame for all landmarks together. Using several landmarks leads to lower mean speeds and thus better motion correction. Also, the mean speed cannot be reduced by additional landmarks once a sufficient number has been reached. This indicates some residual local motion which cannot be compensated by global motion correction.

Table 2: Mean speeds of image sequence, motion globally corrected by mean of increasing number of landmarks, C_5.

Number of landmarks	mean speed
1	2.15 px/frame
2	2.07 px/frame
3	1.99 px/frame
4	2.09 px/frame
5	2.04 px/frame
6	1.97 px/frame
7	1.92 px/frame
8	1.91 px/frame

Finally, five landmarks tracking the motion of the mechanical stabiliser (Fig. 14) are used to correct the motion of the heart surface area stabilised by it. The resulting stabilised image sequence yields a mean speed of 3.57 px/frame. Compared to results achieved by landmarks on the heart surface (Fig. 2) this shows that the motion of the mechanical stabiliser is not sufficiently characteristic of the motion of the heart surface to allow for good motion correction.

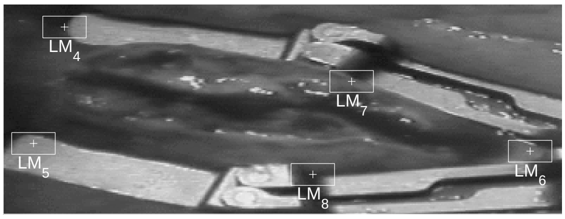

Figure 14: Landmarks at the mechanical stabiliser.

4 RESULTS

Results of motion correction are evaluated by calculating the remaining optical flow in the stabilised images. An analysis of different methods to compute the optical flow shows that a method based on block matching as used for tracking natural landmarks is best suited to estimate the motion field of the beating heart surface reliably. Specular reflections can disturb motion fields considerably but can be eliminated by the approach of (Gröger et al., 2005) before the optical flow is calculated. Moreover, outliers still remaining in the motion field are detected by homogeneity measures and removed before evaluation. The proposed mean speed of image velocities in the motion field proves as a suitable measure to estimate the motion of the heart surface in the image plane and to compare different motion correction strategies.

The evaluation of the motion fields of stabilised images shows that the global motion correction approach is able to reduce the image motion of the beating heart surface significantly. It becomes especially robust when using the mean motion of several landmarks for motion correction. The resulting motion correction consists of a global translation of the image, which can be implemented very efficiently.

The results of global motion correction also show that the motion of the elastic surface of the mechanically stabilised beating heart still varies locally. Therefore image motion cannot be fully reduced by globally constant translation. This global approach corresponds to motion compensation achievable by moving the viewing camera, which is therefore not sufficient either to fully stabilise beating heart images.

5 CONCLUSION

Investigations show that the developed optical flow strategy is able to estimate the motion on the beating heart surface reliably. The proposed measure of mean speeds proves well suited to analyse and compare motion fields describing the motion of the beating heart surface. The presented global motion correction scheme is applicable in realtime and significantly reduces the remaining motion of mechanically stabilised images of the beating heart. However, residual image motion still remains due to local motion differences of the beating heart surface. Therefore, for full motion compensation, locally adaptive motion correction strategies have to be applied.

ACKNOWLEDGEMENTS

This work has been supported in part by the German Research Foundation (DFG). The authors also thank U. Haberhauer for assistance with motion analysis and Dr. D.H. Böhm from the Department of Cardiovascular Surgery at the University Hospital of Hamburg-Eppendorf (UKE), Germany, for clinical support and video sequences of the beating heart.

REFERENCES

Barron, J., Fleet, D., and Beauchemin, S. (1994). Performance of optical flow. *International Journal of Computer Vision*, 12:1:43–77.

Ginhoux, R., Gangloff, J., de Mathelin, M., and et al (2004). Beating heart tracking in robotic surgery using 500 Hz visual servoing, model predictive control and an adaptive observer. In *IEEE ICRA*, pages 274–279, New Orleans, USA.

Gröger, M., Kübler, B., and Hirzinger, G. (2004). Selection of qualified colour markers for heart surface tracking. In *BMT*, volume 49-2, pages 192–193, Illmenau.

Gröger, M., Ortmaier, T., Sepp, W., and Hirzinger, G. (2002). Tracking local motion on the beating heart. In *Medical Imaging 2002*, volume 4681 of *SPIE*, pages 233–241, San Diego, USA.

Gröger, M., Sepp, W., and Hirzinger, G. (2005). Structure driven substitution of specular reflections for realtime heart surface tracking. In *IEEE International Conference on Image Processing (ICIP)*, volume 2, pages 1066–1069, Genova, Italy.

Horn, B. (1986). *Robot vision*. MIT Press, Cambridge, Massachusetts.

Horn, B. and Schunck, B. (1981). Determining optical flow. *Artificial Intelligence*, 17:185–204.

Jacobs, S., Holzhey, D., Kiaii, B., Onnasch, J., Walther, T., Mohr, F., and Falk, V. (2003). Limitations for manual and telemanipulator-assisted motion tracking – implications for endoscopic beating-heart surgery. *Ann Thorac Surg*, 76:2029–2035.

Lucas, B. and Kanade, T. (1981). An iterative image registration technique with an application to stereo vision. In *Proc 7th Intl Conf on Artificial Intelligence*, pages 674–679.

Nakamura, Y., Kishi, K., and Kawakami, H. (2001). Heartbeat synchronization for robotic cardiac surgery. In *IEEE ICRA*, pages 2014–2019, Seoul, Korea.