Automatic quantitative analysis of cardiac MR perfusion images

Marcel Breeuwer[1], Luuk Spreeuwers[2] and Marcel Quist[1]

[1]Philips Medical Systems, EasyVision Advanced Development, Best, The Netherlands
[2]University Medical Center, Image Sciences Institute, Utrecht, The Netherlands

ABSTRACT

Magnetic Resonance Imaging (MRI) is a powerful technique for imaging cardiovascular diseases. The introduction of cardiovascular MRI into clinical practice is however hampered by the lack of efficient and accurate image analysis methods. This paper focuses on the evaluation of blood perfusion in the myocardium (the heart muscle) from MR images, using contrast-enhanced ECG-triggered MRI. We have developed an automatic quantitative analysis method, which works as follows. First, image registration is used to compensate for translation and rotation of the myocardium over time. Next, the boundaries of the myocardium are detected and for each position within the myocardium a time-intensity profile is constructed. The time interval during which the contrast agent passes for the first time through the left ventricle and the myocardium is detected and various parameters are measured from the time-intensity profiles in this interval. The measured parameters are visualised as colour overlays on the original images. Analysis results are stored, so that they can later on be compared for different stress levels of the heart. The method is described in detail in this paper and preliminary validation results are presented.

Keywords: cardiovascular MRI, myocardial perfusion, registration, segmentation, quantitative analysis, automation

1. INTRODUCTION

Cardiovascular diseases have become one of the major death causes in the western society. Although our average life time has significantly increased over the last decades, still about one out of five people dies before the age of 65. Cardiovascular diseases are responsible for about 30% of this mortality (figures for Europe). Since the population is aging, it is expected that the number of people suffering from a cardiovascular disease will increase in the coming decades. There is therefore an increased need for efficient and reliable imaging and image-processing methods for the diagnosis and monitoring of these diseases.

Traditionally, echocardiography, nuclear medicine and x-ray angiography are used to image the heart. Recently, Magnetic Resonance Imaging (MRI) has proven to be also a powerful cardiovascular imaging technique. The amount of anatomical detail that can be acquired with MRI is substantially larger than what can be obtained with the traditional imaging techniques. In order to maximally benefit from this increase in anatomical detail, efficient and reliable image-processing tools are needed. The lack of these tools is presently hampering the introduction of cardiovascular MRI into clinical practice.

This paper focuses on the analysis of images from one of the most important cardiovascular MRI examinations: the evaluation of the perfusion of blood in the myocardium (the heart muscle). An automatic quantitative analysis method is described and preliminary validation results are presented. The remainder of the paper consists of the following sections. In Section 2, the basic principle of the evaluation of myocardial perfusion with MRI is described and requirements with respect to the processing of myocardial perfusion images are formulated. Then in Section 3 the perfusion analysis method that we have developed is described in detail and in Section 4 preliminary validation results are presented. Finally, in Section 5 our experiences so far are summarized.

2. THE EVALUATION OF MYOCARDIAL PERFUSION WITH MRI

2.1 Myocardial perfusion imaging

MR-based myocardial perfusion imaging is performed as follows. During a period of 30-60 seconds, 1-6 short-axis slices are acquired at different positions through the myocardium (128^2 or 256^2 pixels/image). The patient's electrocardiogram (ECG)

1. Marcel.Breeuwer@philips.com; Marcel.Quist@philips.com; phone +31-40-2762456 (Breeuwer) +31-40-2763868 (Quist); fax +31-40-2763816; Philips Medical Systems, EasyVision Advanced Development, Building QV1, P.O. Box 10.000, 5680 DA Best, The Netherlands.
2. Luuk.Spreeuwers@isi.uu.nl; phone +31-30-2508197; fax +31-30-2513399; University Medical Center, Image Sciences Institute, Heidelberglaan 100, 3584 CX Utrecht, The Netherlands.

Medical Imaging 2001: Image Processing, Milan Sonka, Kenneth M. Hanson, Editors,
Proceedings of SPIE Vol. 4322 (2001) © 2001 SPIE · 1605-7422/01/$15.00

is used to trigger MR scanning, typically one set of slices is acquired per 1-2 heart beats at a point of time close to diastole. The result is a 4-dimensional signal $I(x, y, z, t)$, where x and y are integers that indicate the position of a pixel in an individual image, z is an integer that indicates the position of the image along the long axis of the heart (the axis from base to apex, z is usually called the slice number) and t is an integer that indicates the discrete time (i.e. the heart beat) at which the set of slices was acquired.

To ensure that the position of the scanned slices does not vary over time, the patient is fixated on the scanner table and is requested to hold his/her breath during image acquisition. Shortly after scanning has started, a contrast agent is injected into the patient. The arrival of this agent in the myocardium results in an increase of its intensity in the MR images. Insufficiently perfused parts will have a lower and/or delayed intensity increase. Perfusion imaging is usually performed both when the patient's heart is at rest and when it is stressed (using physically or pharmaceutically induced stress)[1,2].

Figure 1 illustrates the scanning process in greater detail. After a period of free breathing the patient is requested to hold his/her breath. During this first breathhold (of about 10 seconds), so-called baseline images are acquired. These images contain the myocardium and the left ventricle without contrast agent and are later on used to correct the measured perfusion parameters for intensity inhomogeneities (see Section 3.6). The patient is then allowed to take several deep breaths after which the breath should again be held. Shortly before the second breathhold, the contrast agent is injected. During the second breathhold, the contrast-uptake images are acquired.

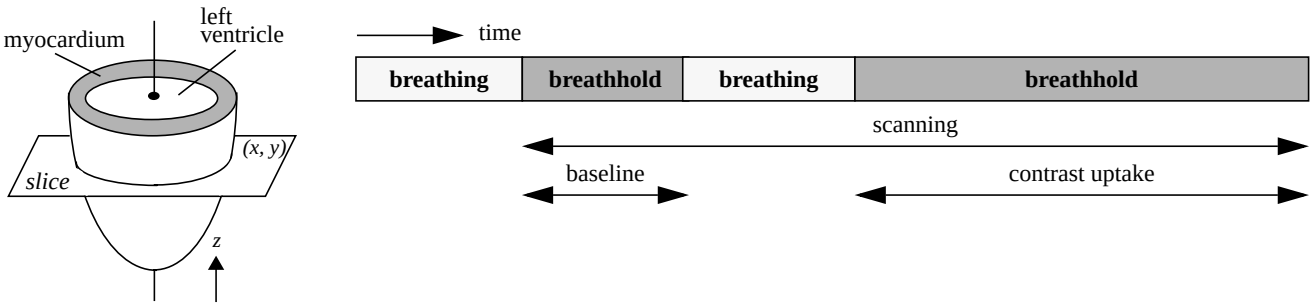

Figure 1: Myocardial perfusion imaging: a) position of a short-axis slice and b) scanning as a function of time

Ideally, the patient holds his/her breath at least until the contrast agent has passed through the myocardium for the first time and the ECG triggering ensures that images are acquired at exactly the same phase of the heart contraction cycle. In this case, the shape and position of the myocardium will be equal in all images and a position (x, y) in an image will unambiguously correspond to a position within the patient. In practice, however, patients often cannot hold their breath sufficiently long, the position of the lungs (and thus of the heart) may differ after the first and second breathhold, and the ECG-trigger equipment may fail due to e.g. the presence of strong magnetic fields. Furthermore, although fixated, the patient may move while lying on the scanner table. As a result, the position of the myocardium may vary over time and so-called 'outlier images' may be present (images scanned at a wrong z position and/or under a wrong orientation). Figure 2a illustrates motion due to breathing and Figure 2b shows one outlier image situated between two correctly scanned images. In this outlier image, the myocardium has a much lower intensity than in the other two images.

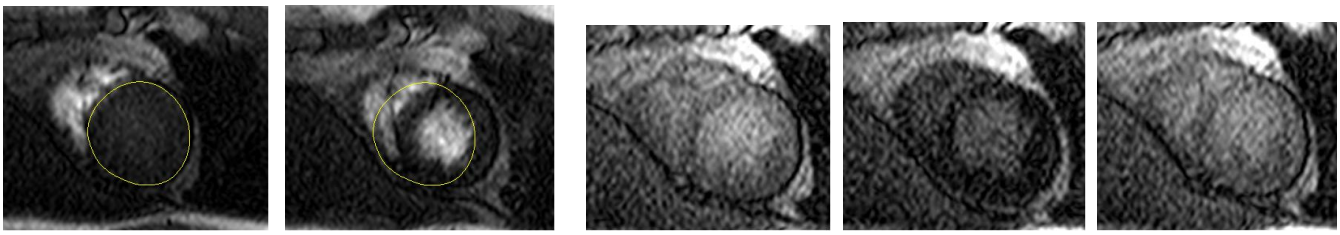

Figure 2: a) Illustration of motion due to breathing, b) example of an outlier image (middle image)

2.2 Quantitative analysis of myocardial perfusion images

The first image-processing step that is needed for the quantitative analysis of myocardial perfusion images is the detection of the myocardium in all images that are involved in the analysis. Due to the motion of the myocardium and the specific properties of the perfusion images (low resolution, low signal-to-noise ratio), this is not a trivial task. Next, corresponding positions within the myocardium have to be identified in all images and the degree of perfusion has to be quantified from the variation of the intensity over time. Finally, the quantitative analysis results must be visualised in an easily understandable way and the results for stress and rest have to be compared quantitatively[1,2]. All processing steps have to be performed preferably automatically or with a minimum of user interaction, so that the time needed for the analysis can be limited to a couple of minutes.

3. THE PROPOSED QUANTITATIVE ANALYSIS METHOD

3.1 Notations

In the remainder the following terminology and notations will be used. The intensity $I(x, y, z, t)$ in the 4D MR image as a function of time *t* for a particular position *(x, y)* at a particular slice position *z* is referred to as the time-intensity profile (TIP) for this slice and this position. For convenience, where possible the index *z* will be omitted, i.e. we will usually consider the processing of one slice at a time. The time-intensity profiles in the myocardium (MC) and inside the left ventricle (LV), i.e. inside the blood pool (see Figure 1a), will be notated as $I_{MC}(x, y, t)$ and $I_{LV}(x, y, t)$, respectively.

3.2 The developed quantitative analysis method

We have developed an automatic quantitative analysis method consisting of the following steps. First, image registration is used to compensate for translation and rotation of the heart over time (see Section 3.3). Next, the boundaries of the myocardium are detected (see Section 3.4). Then for each position within the myocardium a local time-intensity profile is constructed. From the average time-intensity profiles in the left ventricular blood pool and in the myocardium, the time interval during which the contrast agent passes for the first time through the myocardium is detected (see Section 3.5). Various parameters such as the mean and maximum upslope are measured from the local time-intensity profiles in this interval (see Section 3.6) and the measured parameters are visualised as colour overlays on the original images (see Section 3.7). Analysis results are stored, so that they can later on be compared for different levels of stress (see Section 3.8).

3.3 Motion compensation

Due to breathing and/or patient motion, the position of the heart in the 4D image data may vary as a function of time. Ideally, one would like to compensate this motion in *x*, *y* and *z*. In principle, this can be achieved with 3D image registration techniques. However, due to the fact that only a very limited number of slices (different *z* positions) are available, it is difficult if not impossible to compensate motion in the *z* direction (through-plane motion). We have therefore restricted ourselves to the compensation of motion in the (*x, y)*-plane (in-plane motion). Furthermore, we have assumed that the ECG triggering functions sufficiently well, so that the shape of the myocardium does not vary significantly over time The motion to be compensated will then consist of rotation and translation only.

Figure 3a shows a block diagram of our motion-compensation method. The image $I(x, y, t)$ (from now on *z* will be omitted) acquired at discrete time *t* is registered with respect to the already registered image $I_R(x, y, t-1)$ at time *t* - 1 (the subscript *R* is used to denote registered images). Since the motion of the heart usually differs from that of other anatomical structures in the images, the registration is confined to a region of interest (ROI) around the heart. This ROI can either be specified manually by the user or it can be detected fully automatically[3]. An example of an image with a ROI is shown in Figure 3b. It was experimentally found that the ROI must be so large that it contains both the right and the left ventricle for all images involved in the registration process.

The rigid registration itself is performed by iteratively searching for that rotation and translation that maximizes the similarity between the rotated and translated version of the ROI in $I(x, y, t)$ and the ROI in $I_R(x, y, t-1)$. We experimented with various similarity measures: the energy of the difference, the joint entropy, the mutual information, the normalised mutual information[4] and the normalised cross-correlation (the correlation coefficient) of the two ROIs as calculated from their estimated joint probability of pixel values. We found that the correlation coefficient gives the best results (least remaining motion after registration). The normalised mutual information performed second-best in our experiments.

Once the best rotation and translation are found, the complete registered image $I_R(x, y, t)$ is calculated from $I(x, y, t)$ by using trilinear interpolation.

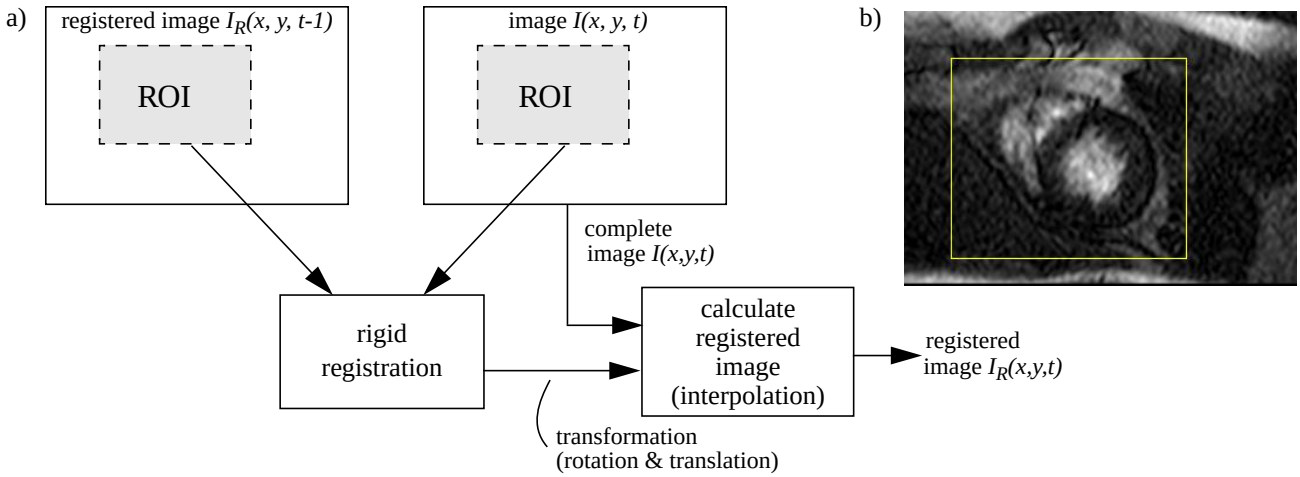

Figure 3: a) Compensation of myocardial motion using image registration, b) example of ROI around the heart

3.4 Contour detection

Once all images have been properly registered, it is relatively easy to visually identify the boundaries of the myocardium. The boundaries can for example be seen very clearly in a temporal maximum intensity projection (MIP) *M(x, y)*, where

$$M(x, y) = \max_t \{I_R(x, y, t)\} .$$

Figure 4a shows an example of a temporal MIP with manually drawn contours. We have however also developed a method of automatic contour detection[3]. Briefly summarized, it operates as follows. First, the centers of the right and left ventricle (RV and LV) are detected. The time-derivative of $I_R(x, y, t)$ is calculated and to suppress noise this derivative is filtered in *x, y* and *t* with a Gaussian filter. To average the perfusion over relatively large areas, the standard deviation of the filter is chosen significantly larger for *x* and *y* than for *t*. The result is a filtered time-derivative image $I_D(x, y, t)$ that has a strong positive peak at the start of the inflow of the contrast agent. The centers of the RV and LV can be easily determined by detecting the positions of the two strongest maxima in a temporal MIP

$$\max_t \{I_D(x, y, t)\}.$$

of $I_D(x, y, t)$. Then, the complete LV and RV are extracted by creating feature images containing for every position *(x, y)* the correlation of $I_R(x, y, t)$ with the time-intensity profiles at the earlier-derived RV and LV centers and by applying straightforward region growing from these RV and LV centers in the feature images.

The perfusion of the myocardium (MC) takes place during and after the passage of the contrast agent through the LV. By subtracting a MIP of images before the passage of the contrast agent from a MIP of images during the maximum perfusion of the MC, a feature image is created that shows a bright LV, a less bright MC and a dark RV and background and thus allows reasonably good separation of MC and background. By using shape constraints (the MC outer boundary is more or less circular) in combination with the previously found LV and RV boundaries, the MC can be found. Figure 4b shows an example of automatically detected contours drawn on the temporal MIP $M(x, y)$ of $I_R(x, y, t)$.

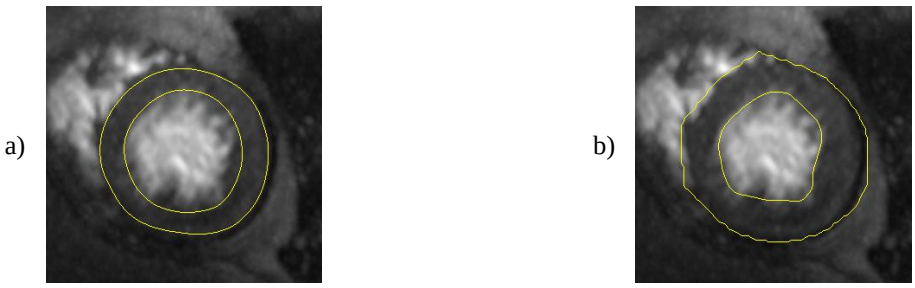

Figure 4: Temporal MIP with a) manually drawn contours, b) automatically determined contours

3.5 Determination of the analysis window

The time window during which the contrast agent passes for the first time through the left ventricle and the myocardium is of major interest for the quantitative analysis[1,2]. This window can be easily determined from the average time-intensity profiles $A_{MC}(t)$ and $A_{LV}(t)$ of the myocardium and the left ventricle, respectively, where

$$A_{MC}(t) = \frac{1}{N_1} \sum_{x,y \in A_1} I_{MC}(x,y,t) \qquad A_{LV}(t) = \frac{1}{N_2} \sum_{x,y \in A_2} I_{LV}(x,y,t)$$

and where A_1 is the set of N_1 pixels inside the myocardium and A_2 is a set of N_2 pixels in the middle of the left ventricle (to avoid inclusion of the papillary muscles, A_2 covers only the central area in the left ventricle). Figure 5 shows typical average profiles and illustrates how the start and end position of the analysis window T_a are determined. First, the time $t_{LV,\,max}$ for which $A_{LV}(t)$ reaches its maximum is determined and the start t_s of the analysis window is defined as that time $t_s < t_{LV,\,max}$ for which $A_{LV}(t)$ has a value lower than or equal to a certain percentage of the maximum value (e.g. 15%). Next, the time $t_{MC,\,max}$ for which $A_{MC}(t)$ reaches its maximum is determined and the end t_e of the analysis window is chosen as $t_{MC,\,max}$ plus a fixed positive offset O. The reason for using this offset is that $t_{MC,\,max}$ only represents the moment at which the average perfusion profile of the myocardium is maximal. Individual positions may reach their maximum at later (or earlier) moments.

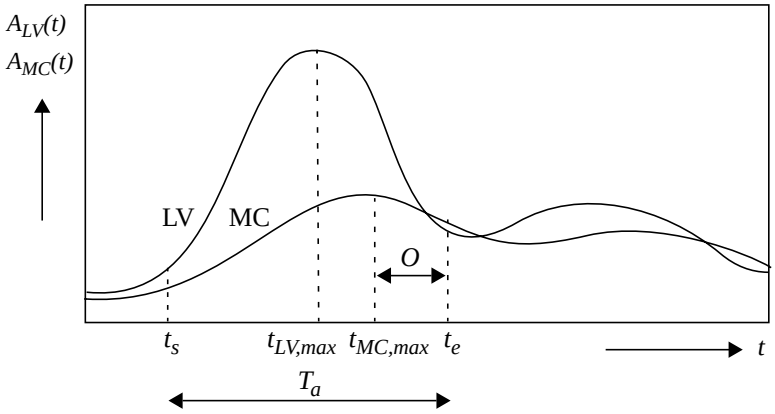

Figure 5: Average time-intensity profile in the MC and LV and the automatically determined analysis window

3.6 Measurement of perfusion parameters

From the time-intensity profile $I_{MC}(x, y, t)$ for $t_s \le t \le t_e$ one or more parameters $p(x, y)$ have to be derived that quantitate the degree of perfusion at position (x, y) in the myocardium. A commonly used parameter is the maximum upslope $a_{MC}(x, y)$ of $I_{MC}(x, y, t)$ (i.e. the maximum value of dI_{MC} / dt)[1,2]. To normalise for the amount of injected contrast agent, the maximum upslope $a_{MC}(x, y)$ is usually divided by the maximum upslope a_{LV} of the time-intensity profile $A_{LV}(t)$. Other possible parameters are the mean upslope within a time window and the time-to-peak t_p (see Figure 6).

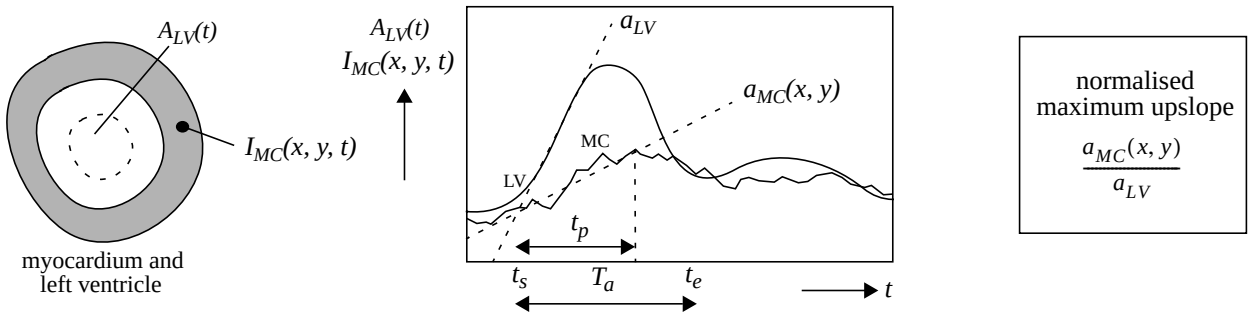

Figure 6: The local time-intensity profile with the calculation of the normalised maximum upslope

Very often, intensity inhomogeneities are present in the scanned images due to the specific properties of the MR receiver coils. These inhomogeneities have to be corrected for, since they significantly influence the time-intensity profiles $I(x, y, t)$

and the parameters derived from this profile. We have assumed that the variation in intensity is a multiplicative effect. We have furthermore assumed that in principle the intensity for different locations in the myocardium should be equal as long as no contrast agent is injected. Under these assumptions, the intensity inhomogeneities can be compensated by dividing the measured intensities *I(x, y, t)* or the calculated upslopes $a_{MC}(x, y)$ by the intensity $\mu(x, y)$ at position *(x, y)* averaged over a number of images in the baseline part of the scan, i.e.

$$\mu(x, y) = \sum_{t=1}^{t_{max}} I_R(x, y, t)$$

where t_{max} indicates a discrete time moment in the baseline scan period (typical value $t_{max} = 5$). A similar correction is performed for the upslope a_{LV}. The corrected upslope is thus calculated as

$$u_{MC}(x, y) = \left(\frac{a_{MC}(x, y)}{\mu(x, y)}\right) \cdot \left(\frac{a_{LV}}{\mu_{LV}}\right)^{-1}$$

The registered images $I_R(x, y, t)$ and associated local perfusion parameters contain a significant amount of noise. It is reasonable to assume that myocardial perfusion will vary relatively smoothly as a function of time and place. To suppress noise in the calculated parameters, the registered images $I_R(x, y, t)$ can therefore be smoothed in place and time (we use a Gaussian filtering with a user-selectable standard deviation).

3.7 Visualisation of perfusion parameters

The locally calculated perfusion parameters are visualised by mapping them to colour values and overlaying the resulting colour map on the original MR images. Furthermore, the myocardium is divided into a number of inner and outer segments, the parameter values are averaged per segment and the result is also displayed as a colour map. The histogram of parameter values inside the myocardium is used to determine which range of parameter values should be mapped to which colours. As is illustrated in Figure 7, the user can select lower and upper percentiles; all values between these percentiles are mapped to the full colour range, all values below/above are clipped.

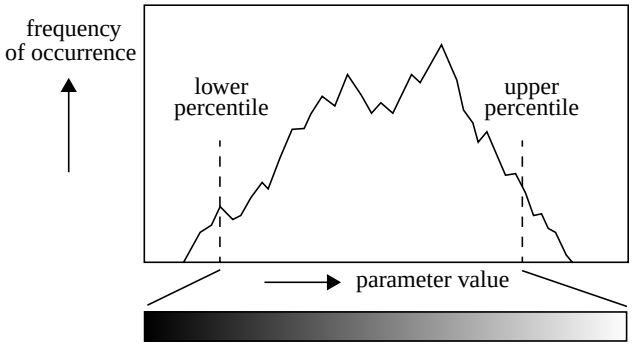

Figure 7: Histogram-based colour mapping of perfusion parameters

Figure 8a shows an example of a visualisation of the locally calculated maximum upslope and Figure 8b shows the corresponding segment-based visualisation (using 12 equally spaced outer and inner segments). Since this paper is printed in black-and-white, a grey scale from totally black to completely white has been chosen instead of a colour scale. Furthermore, the anatomical information outside the myocardium has been omitted. Figure 8c clearly illustrates the effect of the compensation of intensity inhomogeneities. Without this compensation (right image), it could be wrongly concluded that the bottom left part of the myocardium is much better perfused than the other parts.

A very illustrative way of representation of the perfusion parameters are so-called uptake movies or upslope movies. In an uptake movie, the individual images in the time series *I(x, y, t)* are displayed as a movie, while mapping the local intensity of pixels inside the myocardium to colours. For properly registered images, the inflow of the contrast agent can be perceived much better from these colour-coded images than from the original images themselves. In an upslope movie, a time window of user-selectable size is slid over the time-intensity profiles of the myocardium, the maximum or mean upslope is calculated

for every position of the window, and the results are displayed as a movie. Also these movies very clearly show the in- and outflow of the contrast agent in the myocardium. Examples will be demonstrated during the conference.

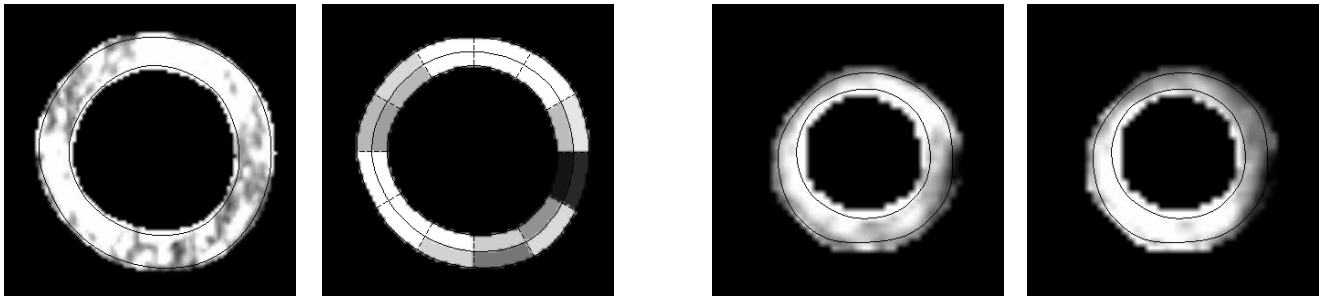

a) b) c)

Figure 8: a) Example of local maximum upslopes, b) corresponding segment-based upslopes, c) example of maximum upslope with compensation of inhomogeneities (left image) and without compensation (right image)

3.8 Comparison of rest and stress

A stenosis in one (or more) of the coronary arteries causes a reduced supply of blood to particular areas in the myocardium. This becomes apparent especially at higher levels of stress. Myocardial perfusion is therefore usually measured at rest and under stress. It has been shown that the perfusion reserve, i.e. the ratio of the maximum upslope under stress and the maximum upslope at rest, is a good indicator of the presence of a coronary-artery disease[1].

In order to measure the perfusion reserve, the myocardium detected from the images that were acquired at rest (the rest scan) has to be registered to the myocardium detected from the images acquired under stress (the stress scan). In general, the shape of the myocardium at rest differs from that under stress, so that a non-rigid registration is needed. We have experimented with two different methods of non-rigid registration, which will be explained in the remainder of this section.

Figure 9a shows the block diagram of a first possible registration approach. Both for the rest and the stress scan a temporal MIP is made from the registered images $I_R(x, y, t)$. The temporal MIPs are then related by means of an affine transformation, which is found by the pixel-based registration technique that was described in Section 3.3. Figure 9b shows an example of the result of this type of registration. The derived affine transformation can be used to relate positions in the myocardium in the rest scan to corresponding positions in the stress scan. It can also be used to derive the position of the myocardial contours in the stress scan from those in the rest scan. Once corresponding positions are known, the perfusion parameters at these positions can be related.

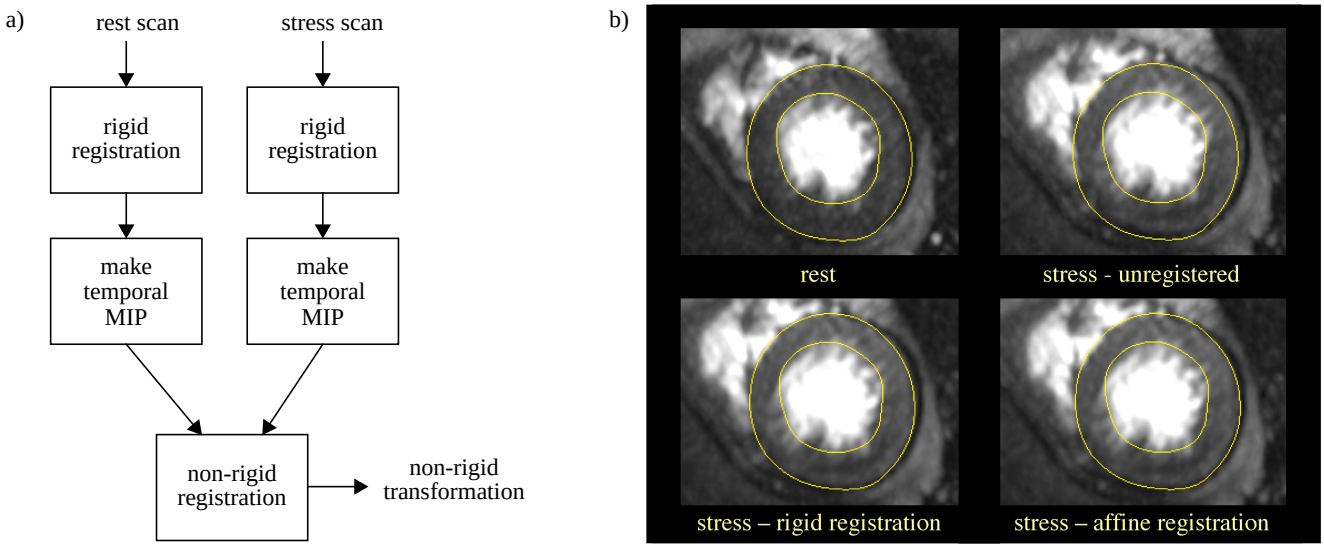

Figure 9: a) Block diagram of non-rigid registration using temporal MIPs, b) example of the result of this registration

Another possible registration approach consists of the non-rigid 'radial' registration of the myocardial contours derived from the rest and stress scans, as is illustrated in Figure 10. First, the centers of the myocardial epicontours are determined. Then, starting from a reference radial, corresponding radials can be compared (e.g. by increasing θ with 1-degree steps). Since the lengths L_r and L_s of corresponding intersections of the radials with the myocardium generally differ, a strategy has to be chosen for relating individual positions on these intersections. We choose straightforward linear mapping. The reference radial can either be specified by the user or it can be automatically derived by means of a pixel-based rigid registration (limited to rotation around the centers only) of e.g. the rest and stress temporal MIPs (including all pixels in a ROI around the heart).

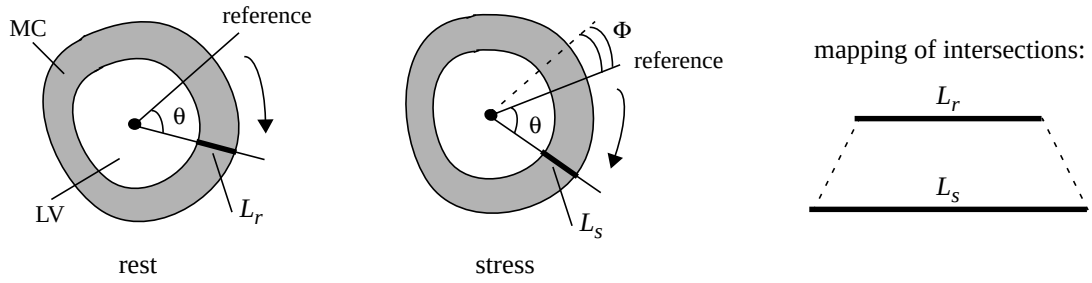

Figure 10: Radial registration of myocardial contours derived from the rest and stress scans

4. PRELIMINARY VALIDATION RESULTS

4.1 Method

The functionality as described in Section 3 was implemented in a prototype perfusion-analysis software package (developed on the EasyScil platform, a Philips Medical Systems proprietary research & development platform comprising functionality of the commercially available EasyVision workstation and functionality of the commercially available SCIL-Image image-processing software package, running on the SUN Ultra workstation family under Solaris 2.6). With this prototype a fully automatic as well as a step-by-step, partially manual analysis can be performed.

For 256 by 256 images, the initial registration (to compensate for myocardial motion) requires in the order of 1 second per image (on a SUN Ultra 10/440 MHz). The automatic ROI detection and the automatic contour detection take in the order of 1 minute each. The complete analysis of one slice requires less than 4 minutes.

The prototype analysis software was evaluated using a total of 30 perfusion scans that were kindly supplied to us by the Cardiology Department of the Deutsches Herzzentrum Berlin (DHZB), Germany and the MR department of the Leeds General Infirmary (LGI), United Kingdom. A cardiologist from the DHZB performed a preliminary clinical validation using scans from 12 patients from whom also x-ray angiograms were available.

4.2 Results and discussion

The results from the preliminary clinical validation are encouraging. For 11 out of the 12 patients scanned at the DHZB from whom also x-ray angiograms were available, a good correspondence was found between the segment-based values of the perfusion reserve as calculated by our prototype and the percentage of coronary-artery stenosis as estimated from the x-ray angiograms.

We found that the initial registration to compensate for myocardial motion is the most crucial part of the analysis. Only if this registration is performed sufficiently well, the automatic contour detection works well and a fully automatic quantitative analysis can be performed. We found that for most scans the registration works well for the contrast-uptake part and for the baseline part, i.e. for the parts during which the patient did not breath heavily. In these parts the motion is relatively small and smooth, so that it can be relatively easily tracked. This is however not the case during the period of breathing between the baseline and uptake parts. Very often, in this period the motion is so large and abrupt that a proper registration cannot be achieved.

As such, the inability to register images scanned during the breathing period is not a problem, since these images are usually not included in the analysis. However, in a significant percentage of the scans, the shape and position of the myocardium after the breathing period differs from that before this period (see the example in Figure 11). A non-rigid registration would be

needed to cope with this problem, which is not yet part of our prototype. This means that for a significant number of scans the intensity-inhomogeneity correction factors (see Section 3.6) could not be determined from the baseline images. In that case, we determined them from a limited number of images at the beginning of the uptake part of the scan, i.e. one or more images in the second breathhold period in which the contrast agent did not yet arrive in the MC. However, not always a sufficient number of images without contrast agent was available, in which case the quantitative analysis could not be performed sufficiently well.

Non-rigid registration may also be useful to improve the motion-compensation process. For example, the size of the myocardium may fluctuate slightly in successive image, which could possibly be corrected for by also allowing scaling in the *(x,y)*-plane.

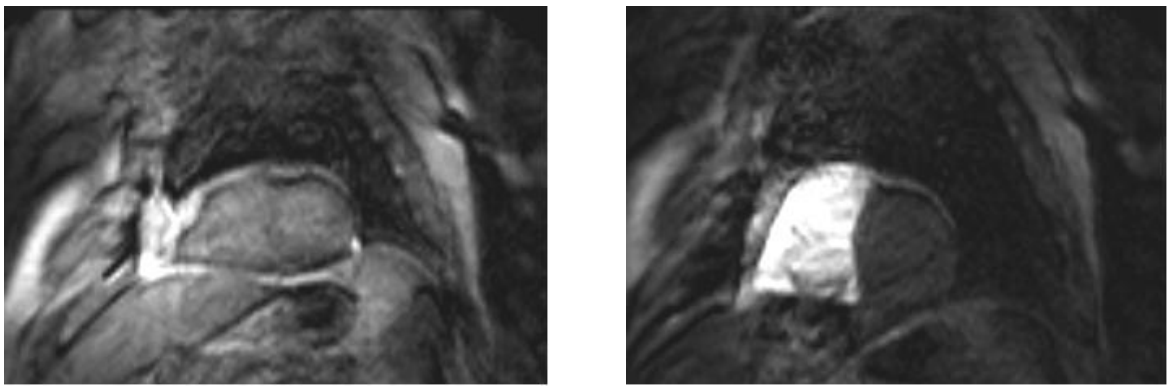

Figure 11: Image directly before (left) and directly after (right) the breathing period

The correction for intensity inhomogeneities is another point of discussion. Even in the case that the baseline images are properly registered to the contrast-uptake images, it can be questioned whether the approach that is currently followed to correct for intensity inhomogeneities is the best one. A problem that has been identified is that in several of the baseline images local intensity variations other than gradients are present in the myocardium (see the example in Figure 12a). At present, these variations are not discriminated from inhomogeneities, which will certainly influence the accuracy of the analysis. Another possible problem is that at the onset of the second scan, some contrast agent may still be present in the blood (see the example in Figure 12b), which will also have its effect on the analysis results. These matters are the subject of current and future study.

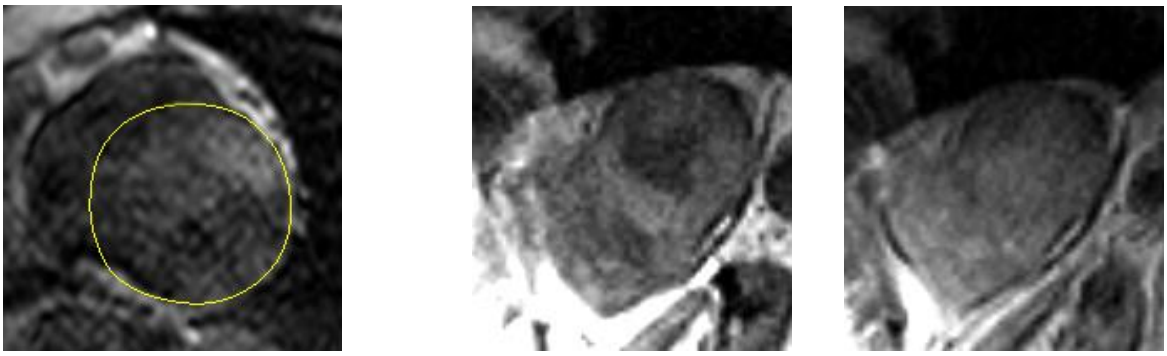

a) b)

Figure 12: a) Intensity variations within the MC and the LV, b) baseline image at rest (left) and at stress (right)

Another subject of ongoing research is the (preferably automatic) detection and removal of outlier slices. Due to failure of the ECG triggering or irregular behaviour of the heart (which is not exceptional for cardiac patients), the position, orientation and/or the time moment at which the image is acquired may significantly vary (see the example in Figure 2b). The detection and removal of outlier images will most probably increase the quality of the motion compensation.

5. CONCLUSIONS

We have introduced a method for fully automatic quantitative analysis of myocardial perfusion images acquired with MRI. Preliminary validation results are encouraging: a good correspondence was found between the values of the perfusion reserve as calculated with our method and the percentage coronary-artery stenosis as estimated from the x-ray angiograms. We are therefore convinced that our approach is a valuable step towards the introduction of MRI-based quantitative myocardial perfusion analysis into daily clinical practice. The method will be more thoroughly clinically validated in the near future.

ACKNOWLEDGEMENTS

We are grateful to Dr. Ingo Paetsch, Dr. Eike Nagel and Dr. Nidal Al-Saadi of the Cardiac MR Department of the Deutsches Herzzentrum Berlin and to Dr. John Ridgway of the MR Department of the Leeds General Infirmary for kindly supplying us with image material and for the discussions with respect to the requirements for quantitative perfusion analysis. We are especially grateful to Dr. Ingo Paetsch for taking the time to perform a preliminary clinical validation of our prototype perfusion-analysis software on 12 patient scans.

REFERENCES

1. N. Al-Saadi et al., "Noninvasive detection of myocardial ischemia from perfusion reserve based on cardiovascular magnetic resonance", *Circulation* **101**, pp. 1379-1383, 2000.
2. Wilke et al., "Myocardial perfusion reserve: assessment with multisection, quantitative, first pass MR imaging", *Radiology* **204**, pp. 373-384, 1997.
3. Luuk Spreeuwers et al., "Automatic detection of the myocardial boundaries of the right and left ventricle in MR cardio perfusion scans", *Proceedings of SPIE Medical Imaging* **4322**, February 2001, San Diego, USA.
4. C. Studholm et al., "An overlap invariant entropy measure of 3D medical image alignment", *Pattern Recognition* **32**, pp. 71-86. 1999.