May 24, 2019

Quantifying Iron Overload using MRI, Active Contours, and Convolutional Neural Networks

Andrea Sajewski
Duquesne University

Stacey Levine
Duquesne University

Quantifying Iron Overload using MRI, Active Contours, and Convolutional Neural Networks

Andrea Sajewski, Stacey Levine

April 2019

1 Introduction

Patients affected by diseases such as sickle cell disease, thalassemias, and leukemias often receive numerous red blood cell transfusions to treat complications of the disease or to reduce side effects of disease treatment. Each transfusion increases the concentration of iron in the body, but since the only excretion methods for iron are through blood loss and loss of skin and mucosal cells, iron accumulates in and damages organs such as the liver, heart, and spleen [6]. Iron chelation therapy is successful at decreasing iron content in tissues, although overtreatment poses its own toxicity risks [5]. Therefore, total iron body content should be carefully assessed and monitored in order to guide the chelation therapy. Hepatic iron content (HIC) is considered a reliable marker for total body iron content [6].

Previously, liver biopsy was the primary method of HIC measurement. In recent years, magnetic resonance imaging (MRI) has been accepted as a noninvasive method for accurate HIC quantification, eliminating the pain and risks involved with biopsy [6]. This is possible because iron is paramagnetic, causing magnetic field inhomogeneity which accelerates the MRI signal decay and shortens the effective transverse relaxation time (T2*) or increases the transverse relaxation rate (R2*=1/T2*) [5]. MR signal decay can be measured and fit to a monoexponential decay curve with time constant T2*. The greater the iron content, the faster the signal decay and the lower the T2*, so the higher the R2*. In fact, it has been shown that R2* and HIC by biopsy are linearly related [5, 7].

After the completion of an MRI scan for HIC quantification, typical radiologist workflow includes manual image segmentation by hand-drawing regions of interest (ROIs) to exclude artifacts and blood vessels. Whole liver regions provide less variability in average R2* [5], but are difficult to manipulate in order to avoid blood vessel contamination. Blood vessels are segmented by thresholding a T2* histogram, as blood has a higher T2* than liver tissue. Overall, this can be a time-consuming and reviewer-dependent process. To streamline this workflow, we propose a tool which will automate and standardize the liver segmentation and vessel exclusion process. Thus, the purpose of this study is

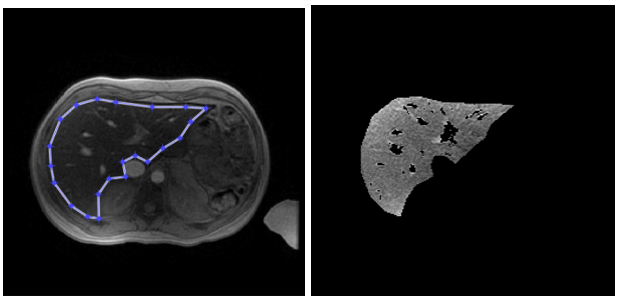

Figure 1: (a) Manual segmentation of 1st image, (b) Vessel exclusion by T2*-thresholding

to investigate the hypothesis that an automated liver segmentation and vessel exclusion method for whole liver ROIs can provide accurate and reliable HIC quantification, improving the effectiveness of therapy and improving patient care.

2 Methods

2.1 Acquisition and T2* Mapping

Anonymized MRI images were acquired from the MeDCaVE dataset. 7 patient scans were used with the following echo times (TE): 5.0, 9.4, 13.7, 18.1, 22.4, 26.8, 31.1, and 35.5 ms.

T2* maps were calculated using custom-written MATLAB software. T2* signal decay was fitted for each pixel with a quadratic mono-exponential model [4], which has been presented recently as an improved fitting approach for accurate and precise T2* quantification in low signal-to-noise conditions and is given by:

$$S^2(TE) = (S_0 e^{-TE/T2*})^2 + N^2. \tag{1}$$

2.2 Manual Segmentation

Whole-liver ROIs were drawn on source images using a MATLAB user interface. A T2* histogram allowed reviewers to exclude pixels with T2* values higher than the main peak and distribution, thus eliminating vessel structures and some pixels of noise from the mean T2* calculation. An overlay mask of the selected and thresholded region was displayed to confirm vessel exclusion (Figure 1).

2.3 Automated Segmentation

2.3.1 Image Denoising: Deep Learning Model

The deep learning technique of convolutional neural networks (CNNs) has been successfully used to classify objects in images and has recently been utilized for image denoising. The Trainable Nonlinear Reaction Diffusion (TNRD) framework [2] is able to denoise images and retain fine features while learning both the convolution filters and nonlinear activation functions. The architecture of this CNN has similar structure to a reaction-diffusion model, where for $t = 1, 2, 3, ...,$

$$\frac{u_t - u_{t-1}}{\Delta t} = -\sum_{i=1}^{N_k} K_i^{t\,\mathsf{T}} \phi_i^t (K_i^t u_{t-1}) - \lambda^t (u_{t-1} - f), \;\; u_0 = f. \tag{2}$$

Here $K_i \in \mathbb{R}^{N \times N}$ is a set of convolution filters, N_k is the number of filters, ϕ_t^i is a learned nonlinear activation function, $\lambda^t > 0$ is a learned a parameter, f is the noisy image, and u_t is the denoised image at time t. The first term in (2) mimics a diffusion term, and the second mimics a reaction term. This model was very effective for pre-processing the MRI data to remove spurious noise before applying the automated segmentation model.

2.3.2 Image Segmentation; Active Contour Model

The active contour model has proven to be an effective tool for image segmentation [1]. This automated model evolves an initial closed contour until it reaches an edge or object boundary, potentially resulting in multiple closed curves that identify ROIs. This model has the potential to isolate the liver from the body MRI, as well as segment out unwanted structures such as blood vessels.

The active contour model evolves a curve, C, to segment an image, $\hat{f}$ based on image edges, according to

$$\frac{\partial C(t)}{\partial t} = g(\hat{f})\kappa\vec{N} - (\nabla g \cdot \vec{N})\vec{N}, \;\; u(t=0) = \hat{f}, \tag{3}$$

where $g(\hat{f})$ is an edge detector that depends on the denoised observed data $\hat{f} = TNRD(f)$, κ is the curvature of C, which controls the speed of evolution, and $\vec{N}$ is the unit normal vector, directed either inward or outward. The first term in (3) controls the speed and direction of the curve evolution, while the second encourages the evolution to stop near likely edges. C is then embedded as a level set of a surface u, which evolves according to

$$\frac{\partial u}{\partial t} = g(\hat{f})|\nabla u|\kappa + \nabla g(\hat{f}) \cdot \nabla u, \;\; u(t=0) = \hat{f}. \tag{4}$$

$\kappa = \text{div}\left(\frac{\nabla u}{|\nabla u|}\right)$ is the curvature of the level lines of image surface u. Embedding C as a level set allows the level line to change topology, that is, it is allowed to 'break' into multiple curves, while evolving according to a continuous

evolution. The final curve is extracted as the zero level set of the solution u of equation (4).

The active contour algorithm is able to isolate the liver from the MRI image after an outward expansion from a user-defined rectangle. The normal vector can then be directed inwards causing the curve to contract and segment blood vessels, which are removed from the region. The active contour algorithm was applied to 3 images for each scan: the first echo (TE = 5.0 ms), the fifth echo (TE = 22.4 ms), and the pixelwise T2* map.

3 Results

We performed several experiments with both manual and automated image segmentations, the results of which are reported in Table 1. First, the average T2* values were calculated in the ROI identified by manual segmentation on the original data, taken to be ground truth. This is reported in the first line of Table 1. We also used TNRD to preprocess the image sequence, performed manual segmentation again, and again calculated the average T2* value which is reported in line 2. We then applied the active contour model to the TNRD denoised image sequence in several ways. First the active contour model was applied to the denoised 1st and then 5th acquired image in the MRI sequence to identify the ROI. Then it was applied to the T2* map generated from the entire denoised image sequence. The average T2* values calculated over these ROIs are reported in lines 3, 4, and 5 of Table 1. The most successful of the fully automated techniques used the T2* map, described explicitly in Algorithm 1 below.

Algorithm 1 Semi-automated T2* Quantification

1. Denoise image sequence and compute T2* map.
   (a) Apply learned TNRD CNN model (2) [3] to denoise each MRI image in sequence (parameters: filter size = 7×7, standard deviation of the noise is $\sigma = 5$)
   (b) Generate T2* map using all denoised images in the sequence
2. Use the T2* map generated from step 1(b) to identify the ROI using the active contour model.
   (a) Prompt user to hand-draw small rectangle on image within the liver
   (b) Apply active contour to image in MATLAB (using the `activecontour` function, Copyright 2012-2015 The MathWorks, Inc.) using rectangle as initial curve
   (c) Using new curve, reverse direction of active contour to segment blood vessels and generate final ROI for the image sequence.
3. Compute the average T2* value within the ROI generated from step 2.

The TNRD CNN was able to denoise the liver MRI images (Figure 2), and the active contour algorithm was, in most cases, able to isolate the liver. Most blood vessels were removed from the ROI. Images displayed are from 3 patient scans (Case 1, Case 3, and Case 4) but are representative of all 7 datasets analyzed (Figure 3, 4, 5). Quantitative results for each of the 7 cases are displayed in Table 1, including all of the calculated T2* values for each method.

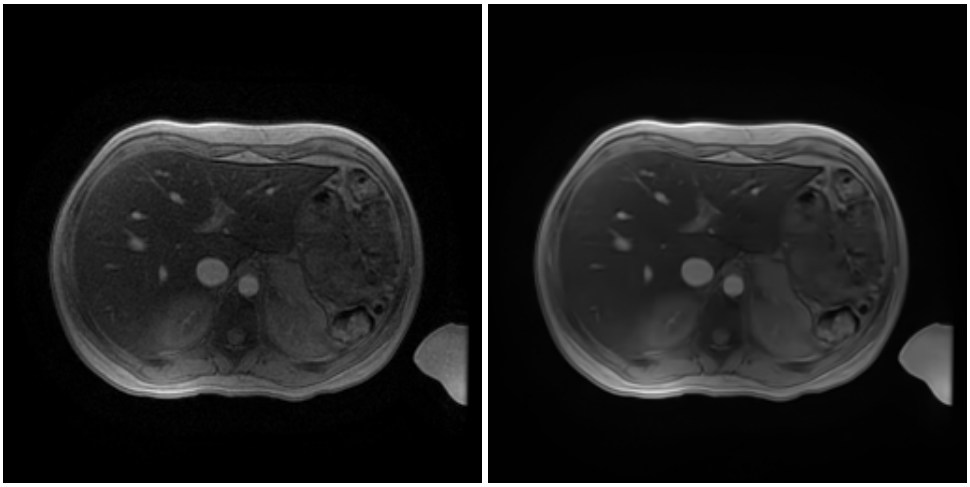

Figure 2: (a) 1st acquired MRI image in Case 4, (b) 1st acquired MRI image, denoised with TNRD

Method	Case 1 T2*	Case 2 T2*	Case 3 T2*	Case 4 T2*	Case 5 T2*	Case 6 T2*	Case 7 T2*
Manual/Ground Truth	9.89	2.58	1279	8.51	1556	29.1	7.92
Manual Denoised	9.82	2.70	1542	8.51	1722	29.4	7.95
1st Image Active Contour	100.5	15.6	1814	128.9	1940	69.2	440.4
5th Image Active Contour	11.0	13.4	1788	22.5	1955	63.2	53.3
T2* Map Active Contour	10.2	3.61	1642	9.09	1756	30.5	27.0

Table 1: Region average T2* results for each segmentation of denoised images

4 Discussion

Manual segmentation produces similar results for denoised images and original images, demonstrating that we do not lose information when denoising. Using the active contour method on the 1st acquired image alone often overestimates T2*, because the image intensity in the liver is still high, so vessels and edges are more difficult to segment. Using the 5th acquired image improves results, especially in cases with low T2* (high iron), but still does not match manual segmentation. Best results were seen from using active contour on the T2* maps generated by the denoised images. Using active contour on the T2* map provides cleaner segmentation and most accurate T2* values as compared to ground truth.

In high T2* (low iron) cases, using the 5th acquired image still overestimates because the signal has not yet decayed enough to see high contrast between liver and surrounding tissue. Thus, some of the higher-intensity tissues around the liver are included in the ROI, and the blood vessels are more difficult to segment out of the ROI. Active contour parameters may be adjusted to account for this, as seen in Figure 5. Parameters that encourage smoothness and direction may need to be optimized based on predicted HIC.

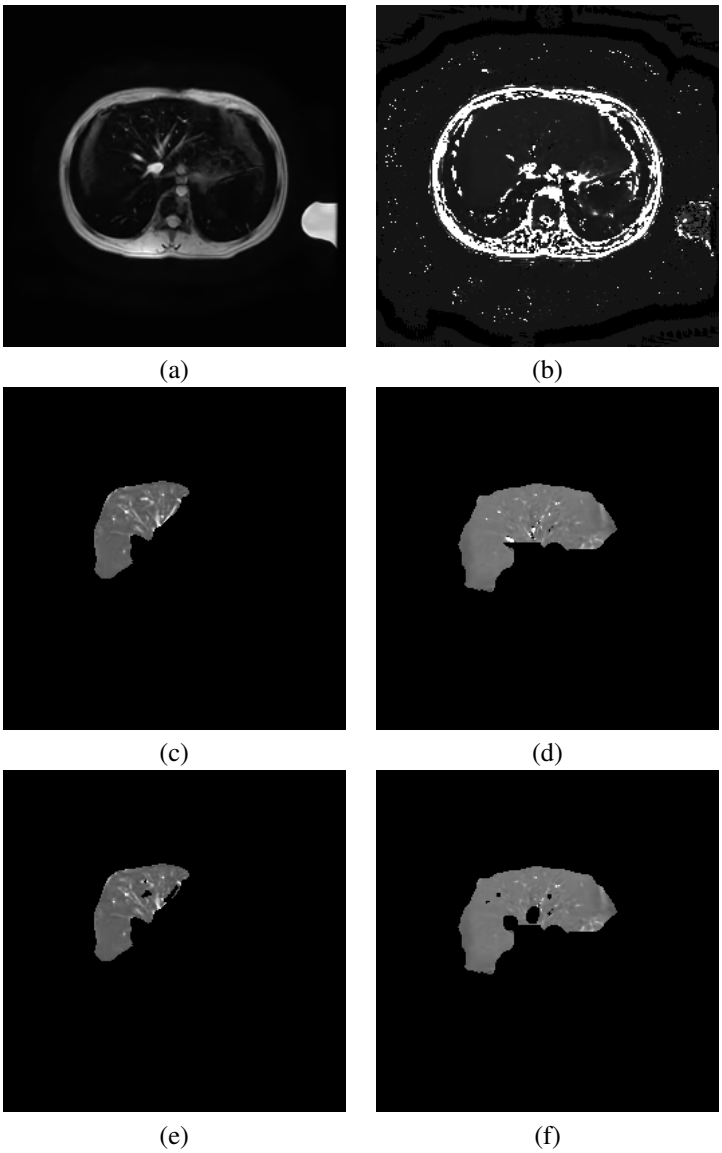

Figure 3: **Case 1:** (a) 5th acquired MRI image, denoised with TNRD. (b) T2* map generated by denoised images. (c) Active contour liver segmentation using denoised 5th image, and (d) T2* map. (e) Active contour exclusion of blood vessels using denoised 5th image and (f) T2* map.

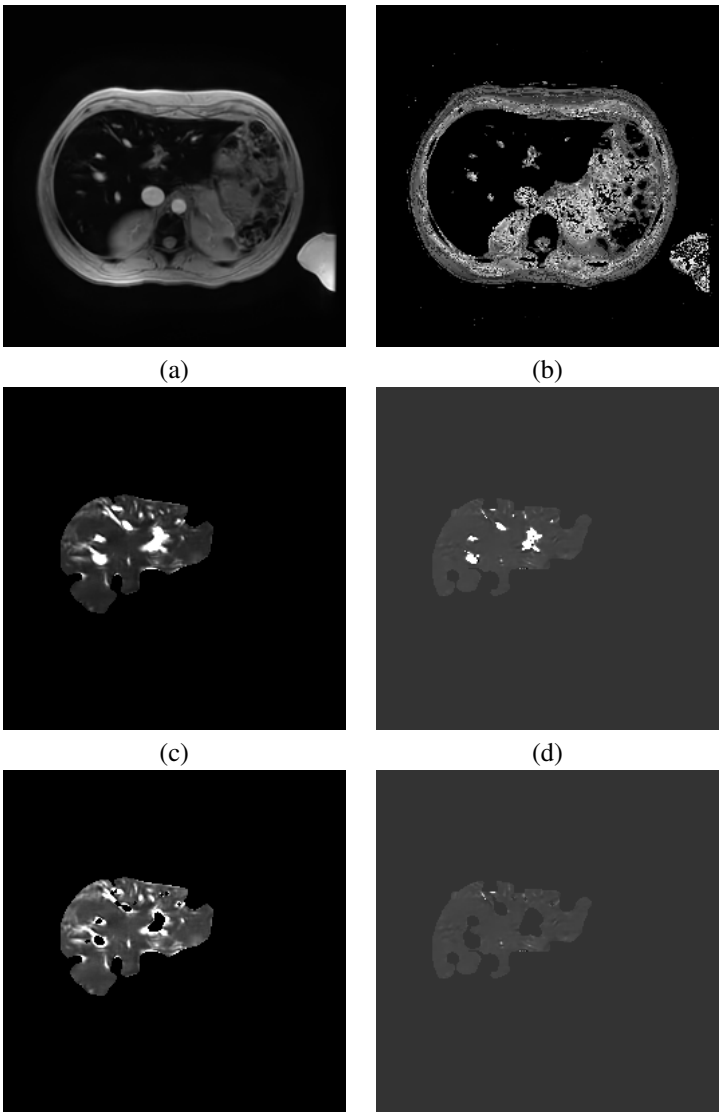

Figure 4: **Case 4:** (a) 5th acquired MRI image, denoised with TNRD. (b) T2* map generated by denoised images. (c) Active contour liver segmentation using denoised 5th image, and (d) T2* map. (e) Active contour exclusion of blood vessels using denoised 5th image and (f) T2* map.

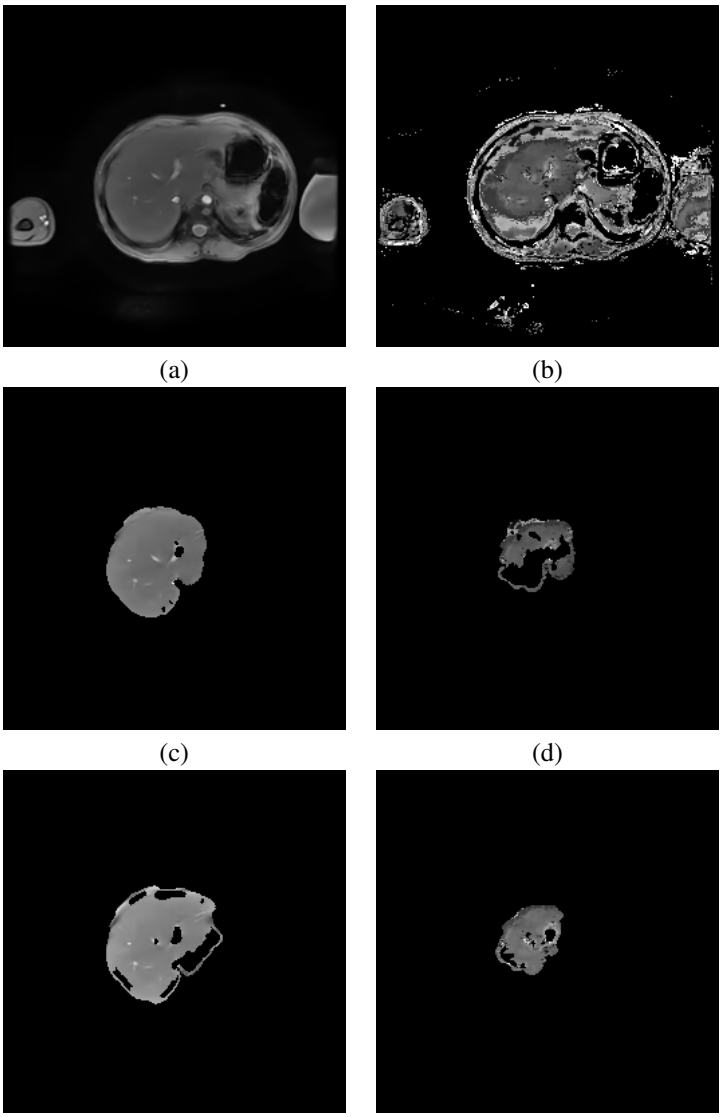

Figure 5: **Case 3:** (a) 5th acquired MRI image, denoised with TNRD. (b) T2* map generated by denoised images. (c) Standard active contour liver segmentation and vessel exclusion using denoised 5th image, and (d) T2* map. (e) Liver segmentation and vessel exclusion using denoised 5th image, adjusting active contour parameters to be less smooth and directed less inward. (f) Liver segmentation and vessel exclusion using T2* map, adjusting active contour parameters to be directed more inward.

5 Conclusion

We conclude that the active contour model is a promising mechanism to isolate the liver from an MRI image and exclude blood vessels in order to calculate T2* and quantify HIC. The active contour model is most promising for accurate segmentation when applied to T2* maps. Segmenting the T2* map using the proposed algorithm provides T2* values comparable to using the current standard of manual segmentation. We found that when using images acquired from the MRI scan, segmentation works better on images acquired later in the sequence. However, any of the experimental methods are useful for predicting high vs low iron content.

We hope to develop an algorithm that incorporates both the image geometry and the expert information from the T2* map in order to more accurately segment the liver MRI. Cross-institutional collaborations could allow for the validation of the active contour algorithm over different imaging protocols, scanners, voxel size, or T2* decay curve fitting methods, and would contribute to an even larger dataset. Ultimately, the goal is to have complete automation of the ROI definition and vessel exclusion process, beginning with standardization of a starting point for the active contour curve, and training the algorithm to optimize parameters to segment the best diagnostic ROI.

This image processing technique could be implemented in MRI scanners and image visualization software at any location, from small clinics to large research institutions. It would benefit small imaging centers who may not have staff onsite with expertise in liver MRI or iron overload who are skilled in segmentation by hand. Using this software would provide them instantaneous, accurate HIC values and diagnoses. Furthermore, this research on the active contour model and TNRD neural network framework for liver and blood vessel segmentation can be applied to other applications in MRI or in other imaging modalities, including segmenting other organs or reducing artifacts on images.

6 Acknowledgements

The authors thank Jonathan Kamis for helping with code and Ryan Cecil for discussion. The authors also thank Drs. Claudia Hillenbrand and Ralf Loeffler for discussion and introduction to the clinical issue.

Anonymized data was used with permission from Dr. Prahlad Menon and the MeDCaVE lab. This study was supported by NSF DMS-1821342.

References

[1] Vicent Caselles, Ron Kimmel, and Guillermo Sapiro. Geodesic active contours. *International Journal of Computer Vision*, 22(1):61–79, Feb 1997.

[2] Yunjin Chen and Thomas Pock. Trainable nonlinear reaction diffusion: A flexible framework for fast and effective image restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(6):1256 – 1272, Aug 2016.

[3] Yunjin Chen, Wei Yu, and Thomas Pock. On learning optimized reaction diffusion processes for effective image restoration. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.

[4] Yanqiu Feng, Taigang He, Peter D. Gatehouse, Xinzhong Li, Mohammed Harith Alam, Dudley J. Pennell, Wufan Chen, and David N. Firmin. Improved MRI R2* relaxometry of iron-loaded liver with noise correction. *Magnetic Resonance in Medicine*, 70(6):1765 – 1774, Jan 2013.

[5] M. Beth McCarville, Claudia M. Hillenbrand, Ralf B. Loeffler, Matthew P. Smeltzer, Ruitan Song, Chin-Shang Li, and Jane S. Hankins. Comparison of whole liver and small region-of-interest measurements of MRI liver R2* in children with iron overload. *Pediatric Radiology*, 40(8):1360–1367, Aug 2010.

[6] Claude B. Sirlin and Scott B. Reeder. Magnetic resonance imaging quantification of liver iron. *Magnetic Resonance Imaging Clinics of North America*, 18(3):359–387, Aug 2010.

[7] John C. Wood, Cathleen Enriquez, Nilesh Ghugre, J. Michael Tyzka, Susan Carson, Marvin D. Nelson, and Thomas D. Coates. MRI R2 and R2* mapping accurately estimates hepatic iron concentration in transfusion-dependent thalassemia and sickle cell disease patients. *Blood*, 106(4):1460–1465, Aug 2005.