

Prior-based Initialization for Automated Analysis of 3D MRE

Bogdan Dzyubak¹, Armando Manduca², Kevin J. Glaser³, and Richard L. Ehman³

¹Physiology and Biomedical Engineering, Mayo Clinic, Rochester, MN, United States, ²Physiology and Biomedical Engineering, Mayo Clinic, Rochester, Minnesota, United States, ³Radiology, Mayo Clinic, Rochester, Minnesota, United States

Purpose: Magnetic Resonance Elastography (MRE) is a clinical technique used for noninvasively diagnosing and staging hepatic fibrosis by calculating shear stiffness from images of mechanical wave propagation in the body [1]. A fully automated technique for the analysis of 2D hepatic MRE has been developed to resolve the problem of inter-reader discrepancies in stiffness measurements from the elastograms [2]. Recently, our implementation of MRE has been extended to image wave propagation in all 3 dimensions allowing for improved accuracy of the stiffness reconstruction. However, this fast EPI-based acquisition produces images with severe intensity inhomogeneity, low edge contrast, and fat suppression (Figure 1) which make them extremely challenging to analyze without manual input as the intensity of a particular tissue type may vary more than the difference between tissues. The method presented here is able to provide reliable initializations for liver segmentation by matching mask priors from a database to the image and cleaning the result based on local and semi-local intensity variations. The target audience for this work is researchers in image processing and quantitative imaging.

Methods: Images: Forty recent 3D MRE exams (32 slices each) with no pathology were retrospectively analyzed from our clinical database in accordance with our IRB protocol. The images were separated into a training set and a test set, and the livers were segmented out manually. The 320 masks from the training set were used as independent liver shape priors while 30 exams (32 slices each) were used as the test set.

Initialization region calculation: The image volumes in the test set were corrected for inhomogeneity with LEMS [3] to reduce intensity variation across the liver (at the expense of making cross-tissue intensities more similar). The initialization area for each exam was then calculated, as follows:

1. The coefficient of variation (CV) image was calculated for each slice (S) by dividing an 8x8 sliding window's means by its standard deviations
2. For each slice (S) of an exam, all images from the training set were rigidly registered to the target slice based on the contour of the inner body wall
3. The prior mask that had the lowest metric s (described below) was selected and the μ and σ of S within the mask were calculated
4. The chosen mask was morphologically dilated, thresholded at $\min(S < \mu - \sigma, S > \mu + \sigma, CV < 0.5)$ and morphologically closed

The metric s was derived by multiplying the prior by S, removing areas with $CV > 0.5$, and dividing the standard deviation of the resulting greyscale mask by the square root of the number of pixels it contained. This metric favors selection of large masks that sample globally homogeneous tissue that may contain some localized outliers that have a high CV.

Evaluation/optimization framework: The automated initialization was compared to manual segmentation in both the test set and the training set, and the fraction of the manually segmented liver captured by the mask (true positive - TP) as well as the fraction of the mask that contained other tissue (false positive - FP) were calculated.

Results: Retroactive matching of prior masks back to the training set using the metric s yielded a TP rate of 90% and FP rate of 9%. Initializations, following the threshold of step 4, yielded a TP and an FP of 75.6% and 10.5% for the training set, and 65.6% and 10.5% for the test set.

Representative best-fit priors and final initialization masks are shown in Figure 2. Of the 1280 images analyzed, 17 had fewer than 100 voxels of liver tissue within the slice. The initialization provided no mask for such slices, which is reflected in the TP rate. The initialization, excluding the homogeneity correction, took approximately 15 seconds per exam.

Discussion: The diagnosis of hepatic fibrosis uses only global measurements from MRE images (typically, the average hepatic stiffness is reported). The method presented here is able to capture nearly 70% of the liver with a low false-positive rate, which is more than sufficient to initialize the segmentation (based on our experience with 2D MRE data), and will allow the processing pipeline to be implemented in a fully automated rather than a semi-automated way.

Although the homogeneity correction makes the liver intensity distribution approximately normal, it exacerbates the problem of intensity overlap between tissues as low-frequency spatial differences are partly attributed to the bias field. The use of intensity priors allows thresholding to be done in a local neighborhood where intensity differences are still present.

The exclusion of areas based on CV prior to the selection of the best fitting mask reduces the impact of hepatic vessels on the selection of the prior. The accuracy of prior fitting back to the training set is not perfect because of this effect. The vascular anatomy is too variable to be accounted for by the priors which, instead, capture the liver outline. Only the largest vessels need to be removed from the liver mask as small vessels, particularly in soft healthy livers, have little effect on the stiffness calculation. Thus, a mask overlap measurement for the quality of vessel exclusion is not meaningful and the vessels are filled in in both manual and automated masks. A short-term improvement for the method would be a better image registration, which is made complicated by the variable amount of fat-suppressed adipose tissue and its broad intensity range. The ability to reliably find a model of the inner body wall in each image to be used for registration would improve the accuracy of prior-matching and reduce the number of priors needed in the database. Additionally, successful whole liver segmentations would be added to the database of priors which would then be reduced to probabilistic atlases of a number of common liver shapes.

Conclusions: The initialization method presented here is able to capture the majority of the liver in EPI-based 3D MRE images that have low contrast, severe intensity inhomogeneity and fat suppression. The method removes the need for manual input and allows a fully automated 3D MRE processing pipeline to be constructed.

References:

- [1] Yin, M., et al., *Assessment of hepatic fibrosis with magnetic resonance elastography*. Clin Gastroenterol Hepatol, 2007. **5**(10): p. 1207-1213.
- [2] Dzyubak, B., et al., *Automated liver stiffness measurements with magnetic resonance elastography*. JMRI, 2013 **38**(2):371-9.
- [3] Salvado, O., et al., *Method to Correct Intensity Inhomogeneity in MR Images for Atherosclerosis Characterization*. IEEE Transactions on Medical Imaging, 2006 **25**(5): p 539-552.

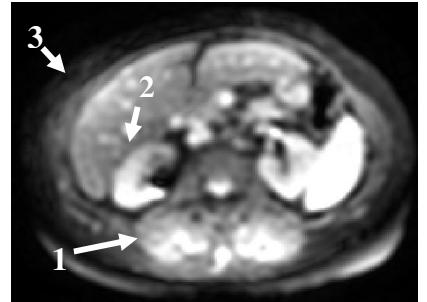


Figure 1: Typical 3D MRE image with inhomogeneous intensity (1), low contrast (2), and fat suppression (3).

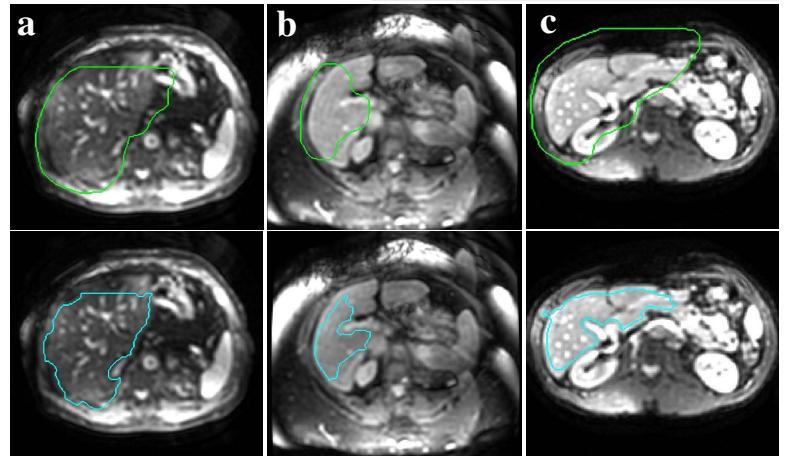


Figure 2: Selected priors (top) and cleaned initialization contours (bottom). The images a-c demonstrate examples of a) a well-fitting prior being used to initialize a low-contrast image, b) a lower lobe of the liver being initialized with a conservative mask, and c) a case where the chosen prior fits poorly due to lack of a good match but the cleanup based on liver statistics combined with local variation filtering results in a high-quality initialization area.