Kidney detection and real-time segmentation in 3D contrast-enhanced ultrasound images

Raphaël Prevost, Benoît Mory, Jean-Michel Corréas, Laurent D. Cohen, Roberto Ardon

KIDNEY DETECTION AND REAL-TIME SEGMENTATION IN 3D CONTRAST-ENHANCED ULTRASOUND IMAGES

Raphael Prevost[1,2], Benoit Mory[1], Jean-Michel Correas[3], Laurent D. Cohen[2], Roberto Ardon[1]

[1] Medisys Research Lab, Philips Healthcare, Suresnes, France
[2] CEREMADE, UMR 7534 CNRS, Paris Dauphine University, Paris, France
[3] Department of Adult Radiology, Necker Hospital, Paris, France

ABSTRACT

In this paper, we present an automatic method to segment the kidney in 3D contrast-enhanced ultrasound (CEUS) images. This modality has lately benefited of an increasing interest for diagnosis and intervention planning, as it allows to visualize blood flow in real-time harmlessly for the patient. Our method is composed of two steps: first, the kidney is automatically localized by a novel robust ellipsoid detector; then, segmentation is obtained through the deformation of this ellipsoid with a model-based approach. To cope with low image quality and strong organ variability induced by pathologies, the algorithm allows the user to refine the result by real-time interactions. Our method has been validated on a representative clinical database.

Index Terms— Kidney, Detection, Segmentation, 3D Ultrasound, Contrast, CEUS

1. INTRODUCTION

Three-dimensional real-time visualization of vascularization can be achieved with CEUS imaging, and provides very useful information for lesions diagnosis or large vessels monitoring [1]. Gas-filled microbubbles, acting as amplifiers of the blood backscattering signal, are used as a contrast agent. Because the bubbles are naturally eliminated by metabolism processes, this modality is considered as completely safe for the patients even with renal or liver failure.

However the usually poor quality of CEUS images makes any analysis challenging: in addition to having powerful speckle noise, the image is very grainy and almost binary as a result of ultrasound interactions with individual bubbles. Unlike in conventional US [2], very few segmentation methods of 3D CEUS images have already been reported.Among them, the authors of [3] propose an interactive method which is specific to tumour segmentation. In [4], an automatic algorithm is developed to segment the heart left ventricle. This method, although applicable to other organs, does not provide any natural way to refine or correct the result interactively.

Our work presents an automatic and fast solution to the problem of kidney segmentation in CEUS images, which has not been tackled yet. While providing a fully automatic result, the proposed solution also allows user interactions with real-time response.

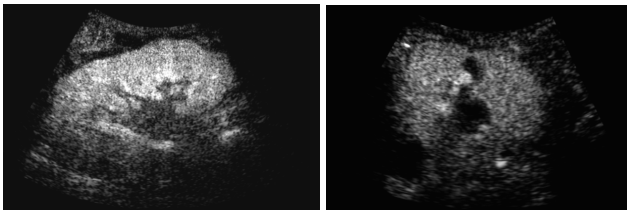

Fig. 1. Slices of two different abdominal 3D CEUS images.

We first detect the kidney by finding a rough estimate of its center, orientation and scale - which is done with a robust ellipsoid detector introduced in Section 2. The outcome of this step is then used as the prior model of a template deformation algorithm described in Section 3.1. It consists in maximizing the image gradient flux through the segmentation border. It is important to use a model-based algorithm: the kidney may not completely lie in the field of view, so the segmentation should be extrapolated outside the ultrasound cone. Furthermore, it may be necessary to guide the segmentation process because of the inherent ambiguities in the images. In Section 3.2, we provide the clinician with a built-in way to correct the result in real-time. A validation of each step of our method is shown in Section 4 while Section 5 concludes this paper and provides material for future work.

2. KIDNEY DETECTION BY ROBUST ELLIPSOID ESTIMATION

Dynamic CEUS images of a kidney show a cortical enhancement shortly followed by a medullary enhancement. Our working images (see Figure 1) are acquired a few seconds after the contrast agent injection. Better visualization of kidney tissue is then available as it is highly hyperechoic whereas its fatty surrounding produces no signal. Since kidney shape can be roughly approximated by an ellipsoid, the kidney detection problem in CEUS images can be initially reduced

to finding the smallest ellipsoid encompassing most of the hyperechoic voxels. Methods such as Hough transforms (e.g. [5]) have already been proposed to detect ellipses in images. However their extension to 3D, though possible, are usually computationally expensive mainly because of the number of parameters to estimate (9 for a 3D ellipsoid). On the other hand, statistical approaches like robust Minimum Volume Ellipsoid (MVE) estimators [6] are better suited but require prior knowledge on the proportion of outliers (here the noise and artifacts), which may vary from one image to another and is thus not available.

We propose a novel algorithm to robustly estimate the ellipsoid's center $\mathbf{c} \in \mathbb{R}^3$ and size/orientation encoded by a 3×3 positive-definite matrix $\mathcal{M}$. Robustly excluding outliers is done by estimating a weighting function w (defined over the image domain Ω into $[0, 1]$) that provides a confidence score for any point $\mathbf{x}$ to be an inlier. Let $I : \Omega \subset \mathbb{R}^3 \to \mathbb{R}^+$ be the grayscale volume, we search $\mathbf{c}$, $\mathcal{M}$ and w as minimizers of the following detection energy:

$$E_d(\mathbf{c}, \mathcal{M}, w) = -\int_\Omega \phi(\mathbf{x})\, I(\mathbf{x})\, w(\mathbf{x})\, d\mathbf{x} + \mu . \log\left(\frac{\mathcal{V}ol(\mathcal{M})}{|\Omega|}\right) . \left(\int_\Omega I(\mathbf{x})\, w(\mathbf{x})\, d\mathbf{x}\right) \quad (1)$$

$$\text{with} \quad \phi(\mathbf{x}) = 1 - (\mathbf{x}-\mathbf{c})^T \mathcal{M} (\mathbf{x}-\mathbf{c})$$

$$\text{and} \quad \mathcal{V}ol(\mathcal{M}) = \frac{4\pi}{3}\sqrt{\det \mathcal{M}^{-1}} \quad \text{the ellipsoid volume.}$$

The ellipsoid is implicitly represented by ϕ (which is positive inside), thus the first term of E_d induces the ellipsoid to include as many bright voxels as possible. The role of w is to neglect the influence of outliers. The second term penalizes the volume of the ellipsoid $\mathcal{V}ol(\mathcal{M})$ with respect to the domain volume $|\Omega|$. It is weighted by a trade-off parameter $\mu > 0$ and normalized by $\int Iw$.

E_d has a statistical meaning: when w is fixed, its minimizers $(\mathbf{c}^*, \mathcal{M}^*)$ are respectively the centroid and proportional to the inverse of the covariance matrix of all voxels, weighted by Iw. Besides, E_d is linear with respect to w which is by definition restricted to $[0; 1]$. Therefore, at every voxel $\mathbf{x}$ the minimizer $w^*(\mathbf{x})$ is equal to 0 or 1, depending only on the sign of $\phi - \mu \log\left(\frac{\mathcal{V}ol(\mathcal{M})}{|\Omega|}\right)$. w^* is then the indicator of the current ellipsoid estimation which has been dilated proportionately to μ.

The choice of μ is paramount. For an ideal case (white ellipsoid on a black background), the algorithm provides the exact solution if $\mu = \frac{1}{4}$ (in 2D) or $\mu = \frac{1}{5}$ (in 3D). In practice, values close to these ones should be chosen.

The minimization of E_d is performed with an alternate iterative scheme that successively updates the variables $\mathbf{c}$, $\mathcal{M}$ and w, as summarized in Algorithm 1. As the energy E_d decreases at each step, the algorithm usually converges. In practice, few iterations are required for convergence and computation time is less than a second on a standard computer. Examples are shown for a synthetic image (Figure 2) and real 3D data (Figure 3).

Algorithm 1: Robust ellipsoid detection algorithm

initialization $\forall\, \mathbf{x} \in \Omega,\ w(\mathbf{x}) \leftarrow 1$

repeat

  // Estimation of center $\mathbf{c}$ and matrix $\mathcal{M}$

  $\mathbf{c} \leftarrow \frac{1}{\int_\Omega Iw} \int_\Omega I(\mathbf{x})\, w(\mathbf{x})\, \mathbf{x}\, d\mathbf{x}$

  $\mathcal{M}^{-1} \leftarrow \frac{2}{\mu \int_\Omega Iw} \int_\Omega I(\mathbf{x})\, w(\mathbf{x})\ (\mathbf{x}-\mathbf{c})\,(\mathbf{x}-\mathbf{c})^T\, d\mathbf{x}$

  // Update of the weighting function w for each $\mathbf{x} \in \Omega$

  if $(\mathbf{x}-\mathbf{c})^T\, \mathcal{M}\, (\mathbf{x}-\mathbf{c}) \leq 1 - \mu \log\left(\frac{\mathcal{V}ol(\mathcal{M})}{|\Omega|}\right)$ **then**

    $w(\mathbf{x}) \leftarrow 1$

  else

    $w(\mathbf{x}) \leftarrow 0$

until *convergence*;

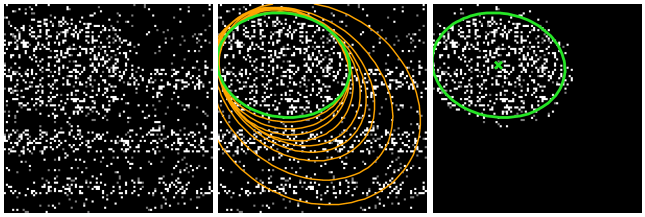

(a) (b) (c)

Fig. 2. (a) Original 2D synthetic image, corrupted by salt-and-pepper noise. (b) Evolution of the ellipse. (c) Ellipse contour and center superimposed on the weighting function wI at convergence.

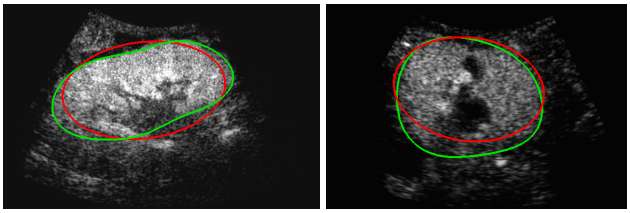

Fig. 3. Results of the ellipsoid detection (red) compared to the ground truth (green), on a slice of the volumes shown in Figure 1.

3. KIDNEY SEGMENTATION

3.1. Segmentation via template deformation

The previously detected ellipsoid will now be deformed to segment the kidney more precisely. We followed the framework described in [7], which is particularly suited to this task as it provides a segmentation algorithm with the possibility to easily interact with the result in real-time. Our goal is to find a transformation $\psi : \Omega \to \Omega$ such that the gradient flux across

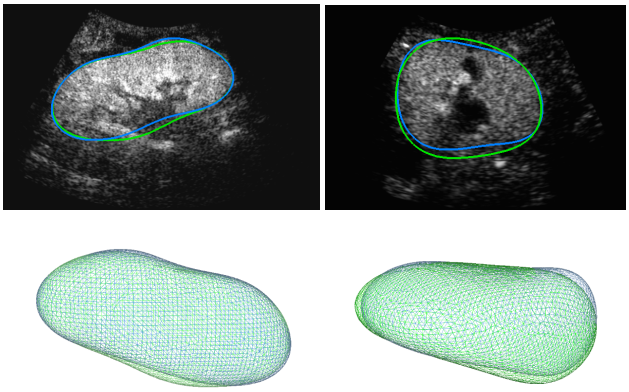

Fig. 4. Result of the automatic segmentation (blue) compared to the ground truth (green) for the two volumes introduced in Figure 1, on a particular slice (top) and in 3D (bottom).

the surface of the deformed ellipsoid $\mathcal{E}(\psi) = (\phi \circ \psi)^{-1}(0)$ is maximum. The segmentation energy is then

$$E_s(\psi) = \int_{\mathcal{E}(\psi)} - \left\langle \vec{\nabla} I(\mathbf{x}) \, , \, \vec{\mathbf{n}}(\mathbf{x}) \right\rangle \, d\mathbf{x} + \lambda \mathcal{R}(\psi) \quad , \quad (2)$$

where $\mathcal{R}(\psi)$ is a regularization term which prevents large deviations from the original ellipsoid. We model the transformation as $\psi = \mathcal{L} \circ \mathcal{G}$ where

- $\mathcal{G}$ is a global transformation, which may correct or adjust the global pose and scale of the ellipsoid (typically a similarity);
- $\mathcal{L}$ is a non-rigid local deformation, expressed using a displacement field $\mathbf{u}$ such that $\mathcal{L}(\mathbf{x}) = \mathbf{x} + (\mathbf{u} * K_\sigma)(\mathbf{x})$. K_σ is a Gaussian kernel that provides built-in smoothness.

This decomposition allows $\mathcal{R}$ to be pose-invariant and constrains only the non-rigid deformation : $\mathcal{R}(\psi) = \mathcal{R}(\mathcal{L}) = \int_\Omega \|\mathcal{L} - Id\|^2 = \int_\Omega \|\mathbf{u} * K_\sigma\|^2$. Finally, using Stokes formula, E_s can be rewritten as

$$E_s(\psi) = - \int_\Omega H(\phi \circ \mathcal{L} \circ \mathcal{G}) \, \Delta I + \lambda \int_\Omega \|\mathbf{u} * K_\sigma\|^2 \quad , \quad (3)$$

where H is the Heaviside function and Δ is the Laplacian operator. This energy is minimized, with respect to the parameters of $\mathcal{G}$ and each component of the vector field $\mathbf{u}$, through a steepest gradient descent.

Some results of this automatic segmentation method are provided in Figure 4 and compared to a ground truth. In more difficult cases, image information is ambiguous due to pathologies and user interactions may be required to correct the segmentation.

3.2. User interactions

The user should be able to indicate that some points $(\mathbf{x}_k)_k$ lie inside or outside the kidney. The implicit representation of the surface allows to formulate these interactions as inequality constraints :

$$\forall \, k \in \{1 \, .. \, n\} \, , \, \gamma_k \, . \, \phi \circ \psi(\mathbf{x}_k) \geq 0 \quad , \quad (4)$$

where $\gamma_k = 1$ (resp. -1) for inside (resp. outside) points. Note that it is also possible to make the kidney surface pass exactly through a specific point by setting it both inside and outside.

To take these constraints into account, we add Lagrangian multipliers $(\alpha_k)_k$ and the energy to minimize is now

$$\tilde{E}_s(\psi) = \max_{\alpha \geq 0} \left\{ E_s(\psi) - \sum_{k=1}^{n} \alpha_k \, . \, \gamma_k \, . \, \phi \circ \psi(\mathbf{x}_k) \right\} \quad . \quad (5)$$

This energy is however not continuous, therefore an augmented Lagrangian optimization scheme is used to solve a smooth approximation of this problem (details are ommitted here but available in [7]). Our current C++ implementation supports up to 100 iterations per second on a standard computer and allows a visualization of the evolution of the segmentation with no latency.

Figure 5 shows an example of a difficult case (presence of cysts, partial visibility of the kidney), for which our method requires some user interactions. With just a few clicks, the user can greatly improve the segmentation result.

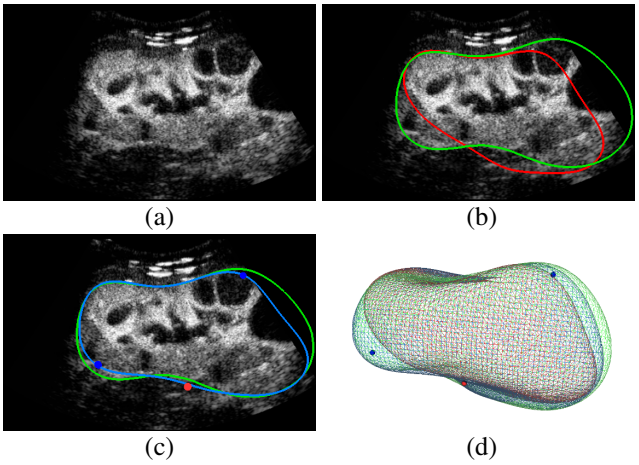

Fig. 5. Example of a segmentation with user interactions. (a) Slice of the original CEUS volume. (b) Comparison of the ground truth (green) and automatic segmentation (red). (c) Corrected segmentation (blue) with 3 clicks. (d) 3D visualization of the ground truth (green), the automatic (red) and corrected (blue) segmentation, with constraint points.

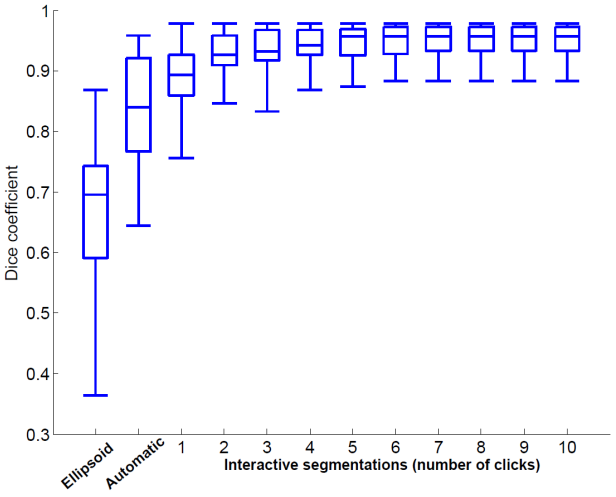

Fig. 6. Boxplot (minimum, lower quartile, median, upper quartile, and maximum over the whole database) of the Dice coefficient between the ground truth and the segmentation at different steps of the proposed algorithm.

4. EXPERIMENTS AND RESULTS

The validation of our method has been performed on a representative clinical dataset of 21 CEUS volumes acquired on a iU22 ultrasound system (Philips, The Netherlands) with different 3D probes (V6-2 and X6-1), spatial resolutions and fields of view. The volumes are composed in average of $512 \times 320 \times 256$ pixels. The 21 patients have been injected with 2.4 mL of Sonovue (Bracco, Italy) contrast agent. For each case, the kidney has been segmented by a radiologist with a semi-automatic tool and is considered as the ground truth.

We compared the results of every step of our method (ellipsoid detection, automatic segmentation, segmentation with interactions) to this ground truth using the Dice coefficient defined as $D(A,B) = \frac{2.|A \cap B|}{|A|+|B|}$. Figure 6 shows the score for the different results as a function of the number of clicks.

The ellipsoid detection produces a median Dice coefficient of 0.696, which is satisfying for an initialization. This score goes up to 0.840 after the automatic segmentation phase. Then, the similarity gradually increases as the user interacts with the algorithm but rapidly converges: most of the time, less than 3 clicks are needed for a fairly precise result. Note that even if the ellipsoid detection provides a bad initialization of the kidney (the minimum score is 0.364), the segmentation may still be successful automatically (≈ 0.644) and with user refinements (> 0.875). The ground truth is not exactly matched because of the high intra-operator variability.

5. CONCLUSION

We proposed in this paper a robust method to segment the kidney in 3D CEUS images. This task is particularly challenging because of the noise, the artifacts and the partial occultation of the organ (due to the limited field of view).

A robust ellipsoid detector has been introduced to coarsely locate the kidney. The ellipsoid is then deformed to segment the kidney more precisely, by maximizing the image gradient flux through the segmentation boundary. We showed that an automatic segmentation may be possible, even if for some difficult and ambiguous images, user interactions are necessary to guide the algorithm. In any case, the kidney segmentation evolves in real-time and only a few seconds are required to obtain an accurate result.

The kidney detection can still be improved by including more anatomical prior knowledge. A possible solution would be to constrain the ellipsoid's axis lengths or volume to be close to clinically meaningful values. Another way is the simultaneous use of CT images of the same patient to extract a tailored model of the kidney and help both the CEUS detection and segmentation. Eventually, we plan to use the kidney segmentation presented here to guide the registration of 3D CEUS with other modalities.

6. REFERENCES

[1] T. Albrecht et al., "Guidelines for the use of contrast agents in ultrasound," *Ultraschall Med*, vol. 25, no. 4, pp. 249–256, 2004.

[2] J. A. Noble and D. Boukerroui, "Ultrasound image segmentation: a survey," *IEEE Transactions on Medical Imaging*, vol. 25, no. 8, pp. 987–1010, July 2006.

[3] A. Gasnier, R. Ardon, C. Ciofolo-Veit, E. Leen, and J.M. Correas, "Assessing tumour vascularity with 3D contrast-enhanced ultrasound: a new semi-automated segmentation framework," in *Proceedings of IEEE ISBI 2010*, 2010, pp. 300–303.

[4] M. Ma, M. Stralen, J. Reiber, J. Bosch, and B. Lelieveldt, "Left ventricle segmentation from contrast enhanced fast rotating ultrasound images using three dimensional active shape models," in *Proceedings of FIMH 2009*, 2009, pp. 295–302.

[5] N. Guil and E.L. Zapata, "Lower order circle and ellipse hough transform," *Pattern Recognition*, vol. 30, no. 10, pp. 1729 – 1744, 1997.

[6] S. Van Aelst and P. Rousseeuw, "Minimum volume ellipsoid," *Wiley Interdisciplinary Reviews: Computational Statistics*, vol. 1, no. 1, pp. 71–82, 2009.

[7] B. Mory, O. Somphone, R. Prevost, and R. Ardon, "Template deformation with user constraints for live 3D interactive surface extraction," in *Proceedings of MICCAI '11 Workshop MeshMed*, 2011.