Published in final edited form as:

Proc IEEE Int Symp Biomed Imaging. 2018 April ; 2018: 635–639. doi:10.1109/ISBI.2018.8363655.

AUTOMATIC BODY LOCALIZATION AND BRAIN VENTRICLE SEGMENTATION IN 3D HIGH FREQUENCY ULTRASOUND IMAGES OF MOUSE EMBRYOS

Jen-wei Kuo¹, Ziming Qiu¹, Orlando Aristizabal^{2,3}, Jonathan Mamou², Daniel H. Turnbull³, Jeffrey Ketterling², and Yao Wang¹

¹Electrical and Computer Engineering, Tandon School of Engineering, New York University, Brooklyn, USA

²F. L. Lizzi Center for Biomedical Engineering, Riverside Research, New York, USA

³Skirball Institute of Biomolecular Medicine, New York University School of Medicine, New York, USA

Abstract

This paper presents a fully automatic segmentation system for whole-body high-frequency ultrasound (HFU) images of mouse embryos that can simultaneously segment the body contour and the brain ventricles (BVs). Our system first locates a region of interest (ROI), which covers the interior of the uterus, by sub-surface analysis. Then, it segments the ROI into BVs, the body, the amniotic fluid, and the uterine wall, using nested graph cut. Simultaneously multilevel thresholding is applied to the whole-body image to propose candidate BV components. These candidates are further truncated by the embryo mask (body+BVs) to refine the BV candidates. Finally, subsets of all candidate BVs are compared with pre-trained spring models describing valid BV structures, to identify true BV components. The system can segment the body accurately in most cases based on visual inspection, and achieves average Dice similarity coefficient of 0.8924 ± 0.043 for the BVs on 36 HFU image volumes.

Keywords

Brain ventricle segmentation; mouse embryo; high-frequency ultrasound; localization; graph cut

1. INTRODUCTION

To study mammalian development, the mouse has been the premier animal model due to its high degree of homology with the human genome. To investigate how genetic mutations manifest themselves during embryonic development as changing in shapes of the brain ventricles (BVs) and other parts of the embryo in 3D views, real-time imaging modalities and fully automatic segmentation algorithms are highly desirable [1].

High-frequency ultrasound (HFU) has become an effective imaging tool for the rapid phenotyping of mouse embryos due to its fast 3D data-acquisition capability and the availability of commercial and research ultrasound scanners [2]. Fig. 1(a) and (b) display selected slices of two HFU image volumes of mouse embryos, one manually truncated to

contain only the head, and another containing the entire scanned volume. As shown in these examples, a loss of ultrasound signal due to either specular reflection or shadowing from overlaying tissues may lead to the loss of boundary contrast and, therefore, BVs and the amniotic fluid regions may contact each other in the HFU images. In order to tackle this issue and leverage the nested structure of different regions in the head region, an automatic segmentation method known as nested graph cut (NGC) has been successfully developed for the segmented HFU images of the mouse embryonic head in a prior work [3]. The NGC works very well in the head images (manually trimmed) because these images contain only nested objects with alternating dark and bright intensities (Fig. 1(a)), but it cannot be applied to the whole-body images effectively because all the objects in the image volume do not form a nested relationship (Fig. 1(b)). Even if we can locate the region inside the uterus, NGC will consider all dark regions in the body part as the BV components because they are all inside the body. Another challenge is that the intensities of BV components can vary greatly within the same image. NGC, which relies on a single threshold to separate dark and bright regions, does not always identify all BV components.

Locating the embryo and BVs from whole-body ultrasound images is a challenging task. To locate the fetus, [4] requires user to select two points manually. [5] automatically initializes the segmentation by registering training images to subject images according to a phase congruency map. This approach works well in BV-focused ultrasound images in [5], but it is very difficult to register whole-body images acquired without focusing on the BVs. Both [4] and [5] applied multiphase level-set-based approach to segment BVs or fetus from 3D ultrasound images according to assumed intensity information and shape priors, but this kind of approach is sensitive to the initialization.

Deep learning has shown great promise for medical image segmentation in recent years. For example, [6] proposes a Hough voting-based convolutional neural networks (CNNs), which can locate midbrain by voting based on results from patch level classifiers. This method has been successfully applied for brain region segmentation in both ultrasound and MRI images of the head. However, this method requires a significant amount of manual segmentation of target regions to train the patch level classifier, which may not be available in practical applications. Furthermore, the trained network may not work well for images of mutants with unseen shapes or intensity distributions in target regions.

In this paper, we propose a novel fully automatic and robust segmentation framework for segmenting both the body and the BVs, shown in Figure 2. First, the input image is pre-processed with an adaptive intensity mapping function and anisotropic filtering to enhance contrast and reduce noise (for details, please refer to [3]). Then a ROI containing the region inside the uterus is located by sub-surface analysis. NGC is then utilized to segment the ROI into four nested parts (BVs, body, fluid outside the body and the uterine wall), from which an embryo mask is generated. In parallel, a multi-level thresholding algorithm is applied to the pre-processed input image to generate initial candidate BVs. These candidate BVs are further masked by the embryo mask to generate additional candidates. Finally, a subset of the candidates that fit pre-trained spring models for BVs are identified. The reason that we apply the embryo mask to the initial candidates is to deal with occasional missing head

boundaries, which can lead to candidates that join true BVs with the fluid region. The main contributions of this work are:

- We develop an algorithm to localize the interior region of the uterus, by analyzing sub-surfaces.
- We develop a method to segment the body region by applying NGC in the uterus region.
- We accomplish BV segmentation by combining the results of NGC and multilevel thresholding with spring model fitting. Multilevel thresholding is introduced to overcome possible intensity distribution variation in the BV regions in the same image. Spring model fitting is robust to the substantial variation in the embryo posture in different images.
- The framework is fully automatic and is robust to the variation in the body posture or shape. It can obtain satisfactory results even when BVs have missing boundaries and inconsistent intensity distributions.

2. METHOD

We describe the major components in Figure 2 in this section.

2.1. ROI Localization

As shown in Fig. 3(a), the whole-body image contains many undesired objects other than four nested objects of interests. In order to apply NGC, we have to locate a ROI containing the interior of uterus. This is obtained by first identifying a set of sub-surfaces (Fig. 3(e)) and then applying a watershed segmentation method to the distance map generated by these valid sub-surfaces.

2.1.1. Sub-surface Extraction—To automatically determine a threshold that can separate dark (corresponding to BV, fluid and cavity) and bright (corresponding to tissue) voxels, we fit a Gaussian mixture model (GMM) with two mixtures to the histogram of the pre-processed image. The threshold is set at the intersection of the two Gaussian distributions. With this estimated threshold, we threshold the pre-processed input HFU image to get surface boundaries between bright and dark voxels (Fig. 3(b)). Afterwards, we remove boundaries with high curvature to keep only those smooth sub-surfaces and identify sub-surfaces through a connectivity analysis (Fig. 3(c)).

2.1.2. ROI Localization by Sub-surface Analysis—There are two steps to locate ROI. The first step estimates the centroid location of the uterus, and identify candidate type 1 sub-surfaces, which refer to those sub-surfaces located between the body and the fluid, and type 2 sub-surfaces, which are located between the fluid and the uterine wall. This is accomplished through an iterative process as follows: (i) Initialize the set of valid sub-surfaces, which include all initial smooth sub-surfaces; (ii) Determine the centroid position of the valid sub-surfaces; (iii) Update the valid sub-surface set based on the direction consistency and distance constraint according to the centroid; (iv) Repeat step (ii)-(iii) until convergence. Fig. 3(d) shows the final valid sub-surfaces and their centroid. The direction

consistency means that a possible type 1 sub-surface's curvature direction and gradient direction should both point to the centroid, and a possible type 2 sub-surface's curvature direction should point to the centroid but its gradient direction should be away from the centroid. We eliminate all sub-surfaces that are neither type 1 nor type 2, and furthermore remove any remaining sub-surfaces that are away from the centroid by more than a threshold, which is the mean plus 1.5 times the standard deviation of the distance between the centroid and sub-surfaces in the valid sub-surface set. The distance between a sub-surface and the centroid is the average of the distances of all points on the sub-surface to the centroid.

As shown in Fig. 3(d), the valid sub-surfaces obtained by the first step can still contain some incorrect sub-surfaces (e.g. some false type 1 sub-surfaces are outside type 2 sub-surfaces). In the second step, we further prune the initial sub-surface set and corresponding ROI iteratively as follows: (i) Sort all valid sub-surfaces according to their distances to the centroid; (ii) Set the initial seed sub-surface set to include the nearest type 1 and type 2 sub-surfaces and obtain the initial ROI by applying watershed segmentation to the distance map generated by seed sub-surfaces; (iii) Get a new ROI by adding the next nearest candidate sub-surface; (iv) check if the new ROI satisfies the constraint below. If yes, add this candidate sub-surface to the seed sub-surface set and remove voxels covered by the new ROI from seed sub-surface set; (v) Repeat step (iii)-(iv) until no additional candidate sub-surface satisfies the checking constraint. Finally, we employ the last updated seed sub-surface set to obtain the final ROI by the watershed algorithm (Fig. 3(e) and 3(f)). The checking constraint is that the new ROI should be larger in volume than the previous ROI and do not contain type 2 sub-surfaces inside. For example, the ROI obtained after adding the top type 1 sub-surface in Figure 3(d) will contain the type 2 sub-surfaces below it, so the top type 1 sub-surface, which is not a true type 1 sub-surface, will not be added to the seed set.

2.2. Segmentation of ROI Using NGC

With the detected ROI that only contains a set of objects satisfying a recursive containment relationship with alternating dark and bright intensities, we now can apply NGC [3] to segment the ROI into four objects (i.e. BVs, body, fluid, and uterine wall). Two segmentation examples are shown in Fig. 4(b) and 4(g), respectively, with blue indicating the body, magenta the candidate BVs, green the amniotic fluid, and yellow the uterine wall. The union of the BV and body regions form the embryo mask. Note that with NGC, the thin boundary of the amniotic sac is sometimes considered part of the body (labeled in blue as in Fig. 4(g)). We apply a morphological opening operation to remove such thin tissues in the fluid region (labeled in green as in Fig. 4(i)).

As shown in Fig. 4(g), some candidate BVs segmented by NGC are incorrect because the detected ROI in this case fails to contain the whole body. In addition, the multiple BVs inside the body often have different intensity distributions. Because NGC uses a single threshold to distinguish between dark and bright region (set as in Sec. 2.1.1), sometimes NGC does not segment all BV regions properly. To overcome both issues, we do not use the BV regions detected by NGC as BV candidates. Rather, we apply multi-level thresholding to generate candidate BVs as explained below.

2.3. Multilevel Thresholding to Generate BV Candidates

In multi-level thresholding, we first select five thresholds based on the GMM distribution obtained in Sec. 2.1.1. Afterwards, with each threshold, we threshold the whole target image and apply morphological processing (closing and opening) to get multiple candidate BVs. Some of these candidates are incorrect because of the missing boundary between BV and fluid (as in Fig. 4(c)). To solve this problem, we use the embryo mask generated by NGC to truncate original candidate BVs to generate additional candidates. All the candidate BVs generated by using different thresholds, and with and without masking, form the final candidate set. Note that we keep the original candidates (Fig. 4(h)) to deal with the occasional case when the generated mask cuts off part of the true body (as in Fig. 4(g)).

2.4. Identify True BV Components Using Spring Models

After NGC and multi-level thresholding, we obtain a set of candidate BVs, some of which are false. To determine the true BV components from candidate BVs, we make use of the prior knowledge about the feasible BV structures. There are usually a main ventricle and two side ventricles. Depending on whether they are connected, there are mainly three types of BV structures [1]. In each valid structure, only 1 to 3 separate connected components should exist and that these components satisfy a certain spatial relationship. We train a spring model to describe each feasible BV structure. We use a shape descriptor to describe each model, which includes the histogram of distances between component voxels and their centroid, the mean and standard deviation of the volumes of components, and distances between each pair of components. The models are determined based on the manually segmented BV volumes. In the experiment, we applied leave-one-out strategy to train the spring models and test the segmentation accuracy.

Given all candidate BVs for an image, we remove those with low probability based on the learned volume distributions to reduce possible candidates. Then we test all combinations of 1 to 3 components among the candidates against each model, and determine the combination with highest fitting score.

3. EXPERIMENTS

3.1. Dataset

The dataset used in the experiments has 36 whole-body HFU images of mouse embryos. All the volumetric ultrasound data are acquired in utero and in vivo from pregnant mice using a 5-element, 40-MHz annular array [2]. The dimension of the 3D images ranges from 150 by 221 by 141 to 210 by 241 by 281 and the voxel size is 50 by 50 by 50 μ m. For all 36 images, manual segmentations of BVs were obtained by a small-animal ultrasound imaging expert.

3.2. Experimental Results and Discussion

Because we only have manual segmentation results for the BVs, we only report quantitative segmentation accuracy for BVs using the Dice similarity coefficient (DSC).

Because there are no other suitable methods that can be directly applied to whole-body ultrasound images for both BV and body segmentation, we compare several variations of our

methods, which differ in generating BV candidates. We consider results that have DSC less than 0.6 as bad results, which are caused by lacking correct candidate BVs. Table I compares the results based on candidates provided by different methods. As shown in Table I, using candidates proposed by NGC only performed worst because of two reasons stated in Sec. 2.2. The segmentation result based on candidates proposed by multilevel thresholding is much better because spring models can combine candidate components from different thresholds to obtain most proper results. However, multilevel thresholding cannot provide correct candidates when the boundary between the true BV and the fluid region is missing (Fig. 4(c)). By using the embryo mask obtained by NGC to truncate candidates proposed by multilevel thresholding, together with the original candidates, we get the most accurate BV segmentation (Fig. 4(d)).

For the body segmentation, visual inspection by the small-animal ultrasound imaging expert has confirmed that most of the results are quite satisfactory. However, mistakes such as the missing head part in Fig. 4(i) still exist due to the deep touching of the body with the uterine wall.

4. CONCLUSION

Segmentation of the body and the BVs in whole-body HFU images is very challenging because of the variety of body posture and shape, BV shape and intensity, and presence of missing head boundaries. The proposed method is fully automatic and robust to such variations. In future work, we plan to further improve the accuracy of body segmentation.

Acknowledgments

The research described in this paper was supported in part by NIH grant EB022950.

REFERENCES

- [1]. Kuo Jen-wei, Wang Yao, Aristizabal Orlando, Turnbull Daniel H., Ketterling Jeffrey, Mamou Jonathan, "Automatic Mouse Embryo Brain Ventricle Segmentation, Gestation Stage Estimation, and Mutant Detection from 3D 40-MHz Ultrasound Data," Ultrasonics Symposium (IUS), 2015 IEEE International pp. 1–4, 2015.
- [2]. Aristizabal Orlando, Mamou Jonathan, Ketterling Jeffrey A. and Turnbull Daniel H., "High-Throughput, High-Frequency 3-D Ultrasound for in Utero Analysis of Embryonic Mouse Brain Development," Ultrasound in Medicine and Biology, vol. 39, no. 12, pp. 2321–2332, 2013. [PubMed: 24035625]
- [3]. Kuo Jen-wei, Mamou Jonathan, Aristizbal Orlando, Zhao Xuan, Ketterling Jeffrey A., and Wang Yao, "Nested Graph Cut for Automatic Segmentation of High-Frequency Ultrasound Images of the Mouse Embryo," IEEE Transactions on Medical Imaging (TMI) vol. 35, no. 2, pp. 427–441, 2015.
- [4]. Dahdouh Sonia, Angelini Elsa D., Gilles Grangé and Isabelle Bloch, "Segmentation of Embryonic and Fetal 3D Ultrasound Images Based on Pixel Intensity Distributions and Shape Priors," Medical Image Analysis, vol. 24, no.1, pp. 255–268, 2015. [PubMed: 25655408]
- [5]. Qiu Wu, Chen Yimin, Kishimoto Jessica, de Ribaupierre Sandrine, Chiu Bernard, Fenster Aaron and Yuan Jing, "Automatic Segmentation Approach to Extracting Neonatal Cerebral Ventricles from 3D Ultrasound Images," Medical Image Analysis, vol. 35, pp. 181–191, 2017. [PubMed: 27428629]

- [6]. Milletari Fausto, Ahmadi Seyed-Ahmad, Kroll Christine, Plate Annika, Rozanski Verena, Maiostre Juliana, Levin Johannes, Dietrich Olaf, Birgit Ertl-Wagner Kai Botzel and Navab Nassir, "Hough-CNN: Deep Learning for Segmentation of Deep Brain Regions in MRI and Ultrasound," Computer Vision and Image Understanding, vol. 164, pp. 92–102, 2017.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Fig. 1:
 (a) The embryonic head HFU image and its segmentation result by NGC [3]. (b) The whole-body image after filtering and its segmentation result. Blue indicates the body, and magenta the BVs.

Fig. 2:
Flow chart of our proposed method.

Fig. 3:

(a) Filtered ultrasound image. (b) Initial sub-surfaces. (c) Remove small sub-surfaces. (d) After step 1 of Sec. 2.1.2, locate centroid and determine valid type 1 (gray) and type 2 (white) sub-surfaces. (e) After the second step of Sec. 2.1.2, obtain seed sub-surfaces. (f) Apply watershed method to the seed sub-surfaces to get ROI.

Fig. 4:

(a)(f) Pre-processed images. (b)(g) NGC Results. (c)(h) Multilevel thresholding results. (d)(i) Final results. (e)(j) Manual BV segmentations. The amniotic membrane falsely segmented as the body in (g) is removed by opening, resulting in final body segmentation (blue) in (i). The DSC values below (b,c,d,g,h,i) show the accuracy of the corresponding candidate BVs (magenta).

Table I:

The mean and standard deviation of DSC for the segmented BVs by different candidate proposing methods (averaged over testing images in cross-validation rounds)

	Average DSC of good results (DSC ≥ 0.6)	Number of bad segmentations among 36 images (DSC < 0.6)
NGC	0.8260 $\pm$ 0.064	12
Multilevel thresholding	0.8865 $\pm$ 0.061	2
Multilevel thresholding + embryo mask	0.8924 $\pm$ 0.043	0