



ELSEVIER

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Feature-based fuzzy connectedness segmentation of ultrasound images with an object completion step

Sylvia Rueda^{a,*}, Caroline L. Knight^{a,b}, Aris T. Papageorghiou^{b,c}, J. Alison Noble^a^a Centre of Excellence in Personalised Healthcare, Institute of Biomedical Engineering, Department of Engineering Science, University of Oxford, Old Road Campus Research Building, Headington, OX3 7DQ Oxford, UK^b Nuffield Department of Obstetrics & Gynaecology, University of Oxford, Oxford, U.K^c Oxford Maternal & Perinatal Health Institute, Green Templeton College, University of Oxford, Oxford, UK

ARTICLE INFO

Article history:

Received 10 June 2013

Revised 28 May 2015

Accepted 11 July 2015

Available online 17 July 2015

Keywords:

Image segmentation

Ultrasound

Shape completion

Fetal imaging

Image quality

ABSTRACT

Medical ultrasound (US) image segmentation and quantification can be challenging due to signal dropouts, missing boundaries, and presence of speckle, which gives images of similar objects quite different appearance. Typically, purely intensity-based methods do not lead to a good segmentation of the structures of interest. Prior work has shown that local phase and feature asymmetry, derived from the monogenic signal, extract structural information from US images. This paper proposes a new US segmentation approach based on the fuzzy connectedness framework. The approach uses local phase and feature asymmetry to define a novel affinity function, which drives the segmentation algorithm, incorporates a shape-based object completion step, and regularises the result by mean curvature flow. To appreciate the accuracy and robustness of the methodology across clinical data of varying appearance and quality, a novel entropy-based quantitative image quality assessment of the different regions of interest is introduced. The new method is applied to 81 US images of the fetal arm acquired at multiple gestational ages, as a means to define a new automated image-based biomarker of fetal nutrition. Quantitative and qualitative evaluation shows that the segmentation method is comparable to manual delineations and robust across image qualities that are typical of clinical practice.

© 2015 The Authors. Published by Elsevier B.V.

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Organ and tissue delineation is essential for underpinning image-based measurements of organ dimensions or tissue region properties. However, manual delineation is a tedious, subjective, time-consuming, and error prone task highly related to the image characteristics and the expertise of the observer. Development of automatic methods for quantitative analysis is especially challenging in ultrasound (US) images, where objects can show strong inhomogeneities and boundaries, can appear fuzzy or are not visible, and in the case of fetal analysis (which motivated this work) further issues are the change in appearance across gestational age and the challenge of fetal movement artefacts. Typically, purely intensity-based methods do not lead to good segmentation results. Several approaches are available at present for segmenting B-mode US images (Noble and Boukerroui, 2006). Among these, the use of local phase, derived from the monogenic signal (Felsberg and Sommer, 2001), has proven useful

for a variety of image analysis tasks including segmentation (Belaid et al., 2011; Hacihaliloglu et al., 2008), registration (Mellor and Brady, 2005), image enhancement (Boukerroui et al., 2001), tissue characterization (Szilágyi et al., 2009), and feature detection (Bridge and Noble, 2015; Mulet-Parada and Noble, 2000; Rahmatullah et al., 2012), since local-phase methods extract structural image information while being invariant to contrast.

Among the many image segmentation methods that are currently available, the fuzzy connectedness (FC) framework can potentially deal with the fuzziness inherently present in US images and is defined by a discrete mathematical formulation, which makes it easy to implement. Fuzzy connectedness (Udupa and Saha, 2003; Udupa and Samarasekera, 1996) is a region-based approach. The main idea consists of defining the strength of local “hanging togetherness” of pixels within an image taking into account their spatial relationship and their intensity similarities within the object of interest. Some variants such as Iterative Relative Fuzzy Connectedness have been shown to be equivalent to other segmentation methods such as graph cuts (Ciesielski et al., 2012) and the Absolute Fuzzy Connectedness with gradient based affinity to level-sets (Ciesielski and Udupa, 2012). This

* Corresponding author. Tel.: +44 1865 617722; fax: +44 1865 617701.
E-mail address: sylvia.rueda@eng.ox.ac.uk (S. Rueda).

approach has proven to be effective in terms of precision, accuracy, and efficiency (Udupa et al., 2006) in segmenting tissues in the presence of intensity gradation in MR and CT images over numerous applications (e.g. Multiple Sclerosis, Udupa et al., 2001, artery-vein separation, Lei et al., 2001, brain tumour segmentation, Moonis et al., 2002, etc.). To our knowledge this article is the first to consider design of a solution specially formulated for US images.

The particular segmentation challenge considered in this paper is 2D fetal US image segmentation. Previous automatic methods developed for this task have focused on extracting standard biometry (size) parameters over a narrow gestational age range. Examples include methods developed for the fetal head (Chalana et al., 1996; Hanna and Youssef, 1997; Lu and Tan, 2000; Lu et al., 2005; Pathak et al., 1997a; 1997b), the fetal femur (Rahmatullah and Besar, 2009; Shrimali et al., 2009; Thomas et al., 1991a; 1991b), and the fetal abdomen (Chalana et al., 1996; Ciurte et al., 2012; Nithya and Madheswaran, 2009; Yu et al., 2008b) by using active contour models, morphological operators, machine learning, deformable models, or Hough transform approaches. Further, there are a limited number of papers in the literature that have proposed to estimate multiple standard fetal biometric measurements using a general method (Carneiro et al., 2008c; Yu et al., 2008a). The latter work, was subsequently translated into a commercial tool, called Auto OB (Carneiro et al., 2008b). Finally, state-of-the-art segmentation methods for automatic biometry of the fetal head and femur were recently compared on ultrasound data acquired across gestational age in a recent medical image analysis challenge (Rueda et al., 2014).

In 3D ultrasound, Yaqub et al. (2014a) has considered segmentation of 3D femur bone volumes using Random Forests. Cuingnet et al. (2013) has considered automatic detection and alignment of the fetal head from 3D US volumes, and automatic standard plane localization from 3D ultrasound volumes has been considered for the fetal abdomen (Ni et al., 2014) and 3D fetal neurosonography (Carneiro et al., 2008a; Yaqub et al., 2014b). Other fetal organs that have been investigated from a quantitative biomedical image analysis perspective are the fetal lungs (Prakash et al., 2002), heart (Deng et al., 2012; Dindoyal et al., 2005; Veronese et al., 2012), fetal face (Feng et al., 2009), and the fetal brain (Namburete and Noble, 2013; Namburete et al., 2012; 2015; Gutiérrez Becker et al., 2010; Yaqub et al., 2013).

Most previous studies were designed to work over a particular gestational age range (particularly 18–22 weeks which corresponds to the interval of the abnormality screening scan). This avoids the main challenges (articulated later in the paper) of developing segmentation solutions applicable across gestation. To our knowledge, the only previous work to propose estimation of a fetal ultrasound biomarker across a large gestational age range is the framework (Namburete et al., 2015; 2014) that accurately predicts the gestational age of the fetus based on analysis of brain structures using a regression forest model.

None of the previous works have attempted to relate the quality of the images to the quality of segmentation results which is an original contribution of this paper, and most prior work only uses a small number of images to develop and validate a method.

As with the work of Namburete et al. (2014) and Namburete et al. (2015) the development of this method was motivated by the clinical need for cost-effective and simple image-based biomarker tools for supporting pregnancy care in the developing world. Ultrasound-based tools are natural to consider for this purpose. Specifically, fetal adipose tissue in the limbs has been shown to be representative of fetal nutritional state (Larciprete et al., 2003), and its quantification has been hypothesised to be a good indicator of fetal growth (Bernstein et al., 1997). Motivated by this, recent clinical studies by our group (Knight et al., 2012a; 2012b) have shown that estimation of adipose tissue from US images of fetal limbs (fat and fat-free regions), via manual delineation, can characterise differences between healthy fetuses and neonates and relates to fetal nutrition. The method proposed in the paper was designed to automate estimation of this image-based biomarker. We are not aware of any previous work on automatic segmentation of arm adipose tissue on fetal US images.

The contributions of this article are three fold. First, we consider how to extend the Absolute Fuzzy Connectedness (AFC) approach to US images by defining a new affinity function. This is done by incorporating information extracted from local phase features instead of image intensities into the AFC framework affinity function. The resulting local phase-based FC framework becomes invariant to contrast and thus is well-suited for US image segmentation. Second, we present a new shape-based method for object completion of one or more ‘gaps’, to deal with missing information resulting from regions without an ultrasonic signal response (for example due to ultrasonic shadows). The result of object completion is then regularised by mean curvature flow. Thirdly, we introduce an approach to quantify the image quality (which can vary considerably between US image acquisitions) of an ultrasound image segmentation validation dataset to appreciate the accuracy and robustness of the developed analysis methodology across clinical data of varying appearance and representativity of potential real world applications. The latter is especially important for US image analysis methods, where results are normally linked to the quality of the images and general practice (with few exceptions) is to report findings on good acoustic window data.

Preliminary versions of parts of this article appeared in Rueda et al. (2011; 2012b). The present paper presents a more general formulation of the complete analysis method, an in-depth evaluation on clinical data, and the new method for quantitative US image quality assessment is introduced for the first time.

The outline of the remainder of the paper is as follows. In Section 2, the overall segmentation framework is introduced and explained in detail. Qualitative and quantitative evaluations, including the proposed method of quantitative image quality assessment, are presented in Section 3. A discussion and conclusions are given in Section 4.

2. Segmentation framework

The overall segmentation framework is composed of several steps summarised in Fig. 1. Each step is explained in the following subsections.

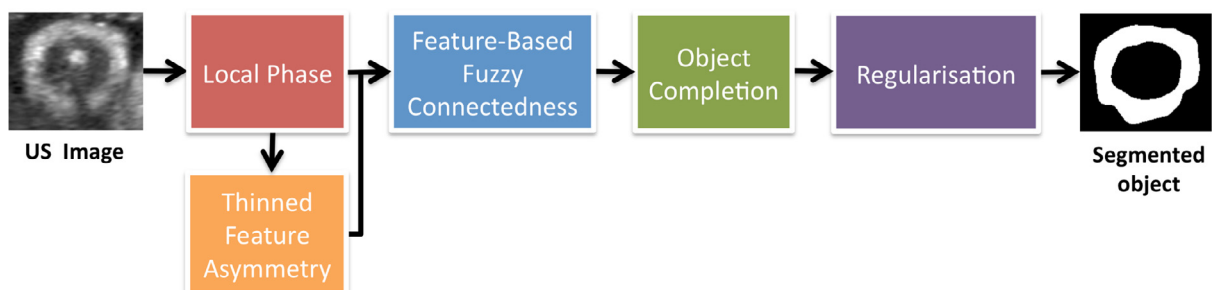


Fig. 1. Proposed feature-based segmentation framework.




Fig. 2. Gaussian derivative kernels for $\sigma > 0$. (a) Even component; (b–c) quadrature pair of odd filters.

2.1. Local phase derived from the monogenic signal

Let $f_A(t)$ be the complex analytic signal derived from $f(t)$ and its Hilbert transform $f_H(t)$ as $f_A(t) = f(t) - if_H(t)$. This representation allows the extraction of the local amplitude (energy) $A(t)$ and local phase $\varphi(t)$ of $f(t)$ defined as $A(t) = \|f_A(t)\| = \sqrt{f^2(t) + f_H^2(t)}$, and $\varphi(t) = \arctan(f_H(t)/f(t))$, respectively.

The monogenic signal (Felsberg and Sommer, 2001) $I_M(x, y)$ of an image $I(x, y)$ generalises the analytic signal to 2D (and higher dimensions) using the Riesz transform instead of the Hilbert transform. From the monogenic signal, the local phase, local energy, and local orientation can be estimated.

In the spatial domain, the convolution kernels of the Riesz transform are defined as

$$h_1(x, y) = \frac{x}{2\pi(x^2 + y^2)^{\frac{3}{2}}} \quad \text{and} \quad h_2(x, y) = \frac{y}{2\pi(x^2 + y^2)^{\frac{3}{2}}}, \quad (1)$$

which in the frequency domain are expressed as

$$H_1(u, v) = \frac{u}{\sqrt{u^2 + v^2}} \quad \text{and} \quad H_2(u, v) = \frac{v}{\sqrt{u^2 + v^2}}, \quad (2)$$

respectively. The quadrature pair (H_1, H_2) define the Riesz transform.

The implementation requires a pair of bandpass quadrature filters to extract the local properties of an image (amplitude, phase, and orientation). The image $I(\mathbf{x})$, where $\mathbf{x} = (x, y)$, is first convolved with a bandpass filter $b(\mathbf{x})$, to give $I_b(\mathbf{x}) = b(\mathbf{x}) \otimes I(\mathbf{x})$, where \otimes denotes the convolution operation. The bandpass filter chosen was a Gaussian derivative filter (Boukerroui et al., 2004) defined in the frequency domain as

$$B(\mathbf{u}) = |\mathbf{u}| \exp(-\mathbf{u}^2 \sigma^2), \quad (3)$$

where $\mathbf{u} = (u, v)$ and σ is the selected scale of the filter.

This filter was empirically chosen, giving better visual maps than other candidate bandpass filters. This is not a critical part of the methodology and other filters, such as Cauchy (Boukerroui et al., 2004), may be better suited for other applications. An example of a Gaussian derivative bandpass filter and the resulting bandpass quadrature pair of odd filters is shown in Fig. 2.

The monogenic signal $I_M(\mathbf{x})$ of $I(\mathbf{x})$ is then expressed as

$$I_M(\mathbf{x}) = (I_b(\mathbf{x}), h_1(\mathbf{x}) \otimes I_b(\mathbf{x}), h_2(\mathbf{x}) \otimes I_b(\mathbf{x})). \quad (4)$$

The local amplitude (energy) $A(\mathbf{x})$, local phase $\varphi(\mathbf{x})$, and local orientation $\theta(\mathbf{x})$ of $I(\mathbf{x})$ are derived from $I_M(\mathbf{x})$ and defined as

$$A(\mathbf{x}) = \sqrt{I_b(\mathbf{x})^2 + (h_1(\mathbf{x}) \otimes I_b(\mathbf{x}))^2 + (h_2(\mathbf{x}) \otimes I_b(\mathbf{x}))^2}, \quad (5)$$

$$\varphi(\mathbf{x}) = \arctan\left(\frac{I_b(\mathbf{x})}{\sqrt{(h_1(\mathbf{x}) \otimes I_b(\mathbf{x}))^2 + (h_2(\mathbf{x}) \otimes I_b(\mathbf{x}))^2}}\right), \quad (6)$$

$$\text{and} \quad \theta(\mathbf{x}) = \arctan\left(\frac{h_2(\mathbf{x}) \otimes I_b(\mathbf{x})}{h_1(\mathbf{x}) \otimes I_b(\mathbf{x})}\right), \quad (7)$$

respectively. The structural information is invariant to contrast and contained in the local phase, whereas the local amplitude represents

the energy, which is dependent on intensity values. An example of local phase image can be seen in Fig. 15(a) for the image in Fig. 11(b).

2.2. Feature asymmetry

Computing the local phase at different scales, allows one to detect step edge features as points where there is local phase congruency (Kovesi, 1999). In other words, a positive step edge will have a local phase value of 0° and a negative step edge will have a value of 180° . To detect step edge features, we use the feature asymmetry FA measure, calculated over a number of scales, and defined as

$$FA(\mathbf{x}) = \frac{1}{N} \sum_s \frac{[|\mathbf{odd}(\mathbf{x})_s| - |\mathbf{even}(\mathbf{x})_s| - T_s]}{\sqrt{\mathbf{even}(\mathbf{x})_s^2 + \mathbf{odd}(\mathbf{x})_s^2 + \varepsilon}}, \quad (8)$$

where $\mathbf{even}(\mathbf{x}) = I_b(\mathbf{x})$, $\mathbf{odd}(\mathbf{x}) = (h_1(\mathbf{x}) \otimes I_b(\mathbf{x}), h_2(\mathbf{x}) \otimes I_b(\mathbf{x}))$, $[\cdot]$ sets to zero the negative values, s represents the scale, N is the total number of scales, ε is a constant that avoids the division by zero (typically $\varepsilon = 0.01$), and T_s is an orientation independent threshold that controls the spurious responses to noise at scale s (Kovesi, 1999; Mulet-Parada and Noble, 2000). T_s can be estimated from statistical properties of the energy response (Kovesi, 1999) or by approximating the statistical mode (Mulet-Parada and Noble, 2000).

The FA image consists of thick detected edges with values close to 1 and with homogeneous regions close to 0 values. An example of a feature asymmetry image can be seen in Fig. 15(b) for the image in Fig. 11(b). However, a good localization of the edges of the object of interest is essential in this framework. Therefore, we need a technique to thin the feature asymmetry edge features while retaining most of the information present in the FA image. A non-maximal suppression technique (such as Sonka et al., 2008) could be used for this, but it will be unable to retain information in directions other than the local orientation direction at each edge pixel. Therefore, a modified non-maximal suppression technique was developed. First, for each pixel in the FA image, non-maximal suppression is performed in all possible directions. Then, at each pixel, the maximum value among all directions is retained. This strategy captures the relevant edge information with good localization while retaining the same intensity present in the FA image, thus obtaining the edge map E that will be used in this work. An example can be seen in Fig. 15(c) for the feature asymmetry image in Fig. 15(b).

2.3. Feature-based fuzzy connectedness

Although several variations of the fuzzy connectedness method exist (e.g. Iterative Relative Fuzzy Connectedness - IRFC, Relative Fuzzy Connectedness - RFC), in this paper we have chosen to employ one of the original formulations of FC, namely Absolute Fuzzy Connectedness (AFC), to study how it would perform using the affinities specially formulated for US imagery.

The Absolute Fuzzy Connectedness strategy (Udupa and Saha, 2003; Udupa and Samarasekera, 1996) is based on a global fuzzy relation that assigns a strength of connectedness to every pair of pixels

in an image to define objects via dynamic programming. The key step of this region-based approach relies on the definition of a *local* fuzzy relation μ_κ , called *affinity*, which defines the local “hanging togetherness” between any two adjacent pixels. If two pixels c and d are adjacent, the affinity depends on how homogeneous the region is and on how close the intensity values at c and d are from the expected intensity value of the object of interest. The affinity is equal to 0 for non-adjacent pixels.

The affinity values are used to define a *global* relation, called Fuzzy Connectedness, where the strength of connectedness between any two pixels is calculated as the largest of the strengths of all paths between c and d on the discrete image grid. Each path corresponds to a sequence of adjacent pixels starting from c and finishing in d and has a corresponding strength value, which is the smallest affinity of any pair of consecutive pixels along the path (weakest link). The Absolute Fuzzy Connectedness is represented as a connectivity map, where the object of interest is obtained by thresholding the image at T_{FC} . The detailed mathematical description of the method can be found in (Udupa and Saha, 2003; Udupa and Samarasekera, 1996).

The initialisation of the general method is based on manually placing one or several seeds within the object of interest. A minimal training stage is required once to define the typical mean and standard deviation of the intensity values of the object of interest.

The AFC framework was adapted to US segmentation by defining a new affinity function that uses structural and edge feature information instead of intensities and intensity gradients. The affinity function was designed as follows. Assume that every fuzzy subset \mathcal{A} in a set is characterised by its membership function $\mu_{\mathcal{A}}$ with values in $[0, 1]$. Given an image, the affinity is composed of three factors: an adjacency component μ_α , an object feature-based component μ_ϕ , and a homogeneity-based component μ_ψ . The adjacency component μ_α is a non-increasing function of the distance in pixels (i.e. integers) $\|c - d\|$ defined as

$$\mu_\alpha(c, d) = \begin{cases} 1, & \text{if } c = d \text{ or } \|c - d\| = 1, \\ 0, & \text{otherwise.} \end{cases} \quad (9)$$

In the original framework by Udupa and Samarasekera (1996), the object feature-based component μ_{ϕ_1} was defined based on the intensities of the image, whereas the homogeneity-based component μ_{ψ_1} was a measure of intensity gradient. The proposed method incorporates the local phase information into the object-feature based component, extracting structural information and making the image invariant to contrast. The edge map E , derived from the feature asymmetry image, directly gives a measure of homogeneity, since smooth regions have small values and regions near boundaries have large values (cf. Section 2.2). Therefore, it is natural to consider it in the definition of the homogeneity-based component. Let $\varphi(c)$ be the local phase at pixel c and $E(c, d)$ the thinned pixel edge derived from feature asymmetry between pixels c and d . The homogeneity-based component μ_{ψ_2} will have a high affinity in homogeneous regions and small affinity at the edges. Since E is close to 0 in homogeneous regions and close to 1 at edge features, we can express the homogeneity component as

$$\mu_{\psi_2}(c, d) = 1 - E(c, d) = g_3(E(c, d)), \quad (10)$$

where g_3 is a function of $E(c, d)$. The object feature-based component μ_{ϕ_2} takes into account characteristic features of the object of interest. In this paper, a recent formulation (Ciesielski and Udupa, 2010) was applied directly to the local phase image instead of intensities, as follows:

$$\mu_{\phi_2}(c, d) = e^{-\max\{\|\varphi(c) - m_o\|, \|\varphi(d) - m_o\|\}^2 / 2\sigma_o^2} = g_4(\varphi(c), \varphi(d)), \quad (11)$$

where m_o and σ_o are the mean and standard deviation of the intensity values of the object of interest, previously calculated in a training stage, and g_4 is a function of $\varphi(c)$ and $\varphi(d)$.




Fig. 3. c -scale estimation at p in a piece of boundary (black curve). $C(p)$ is the c -scale segment associated with p and $C_h(p)$ is the c -scale value corresponding to $C(p)$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article).

There exist several ways of combining the affinity components to form the fuzzy affinity μ_κ (Ciesielski and Udupa, 2010). One general form commonly used is

$$\mu_\kappa(c, d) = \mu_\alpha(c, d)[\omega_1 g_1(I(c), I(d)) + \omega_2 g_2(I(c), I(d))], \quad (12)$$

where $I(c)$ and $I(d)$ correspond to the intensities at pixels c and d , respectively (Udupa and Samarasekera, 1996), $g_1(I(c), I(d)) = \mu_{\psi_1}(c, d)$, and $g_2(I(c), I(d)) = \mu_{\phi_1}(c, d)$. The equivalent affinity function μ_κ^* for the proposed approach is expressed as

$$\mu_\kappa^*(c, d) = \mu_\alpha(c, d)[\omega_1 g_3(E(c, d)) + \omega_2 g_4(\varphi(c), \varphi(d))], \quad (13)$$

where $\omega_1 + \omega_2 = 1$, and with g_3 and g_4 as defined in (10) and (11), respectively.

2.4. Delineating closed regions

The segmentation resulting from the feature-based FC only incorporates regions of the object of interest present in the image. However, it is unable to delineate object boundaries in shadowed areas (e.g. shadow under the humerus bone in Fig. 11(b)), as there is no ultrasonic signal response from these regions. Furthermore, in some cases, the object of interest can be formed by several connected pieces with missing information between them that we would like to retrieve. To overcome this, a new object completion technique has been developed. This first detects the region of the object of interest with missing information, and then fills the gap(s) in by using local shape constraints (Rueda et al., 2008). In the preliminary version of the method (Rueda et al., 2012b), only one gap was corrected. In this paper we have generalised the approach to the detection and completion of any number of gaps appearing in the object of interest after segmentation. The object completion step is described in the next subsection.

2.4.1. c -scale shape descriptor

At each point p on a boundary, a local curvature scale segment (Rueda et al., 2008), called c -scale segment $C(p)$, is defined as the set of connected points at a distance smaller than a threshold t from the line connecting the two end points of the set (red dashed curve in Fig. 3). Each $C(p)$ is obtained after symmetrically and progressively examining the adjacent boundary elements to p until the distance is no greater than a threshold t .

A c -scale value $C_h(p)$ (green dashed line in Fig. 3) can then be obtained as the chord length corresponding to $C(p)$, which is the length of the straight-line segment between the end points in $C(p)$. Large $C_h(p)$ values indicate small curvature at p , whereas small values denote high curvature (Rueda et al., 2008). Values of c -scale are very useful in estimating actual segments and their curvature by considering the local morphometric scale of the object, and are independent of digital effects and noise. The c -scale method has proven to be robust in obtaining a complete description of shape directly applied to digital boundaries. More details can be found in Rueda et al. (2008).

An extension of this implementation was developed to obtain the normal \mathbf{n}_p at each point p in the boundary as the line perpendicular to $C_h(p)$ passing through p (Fig. 3). The direction of the normals is always selected pointing to the inside of the object. Only the normal information is needed for the object completion step, which is described in the following.

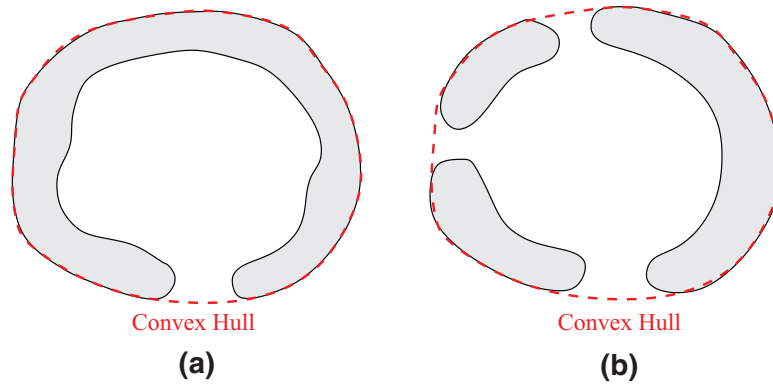


Fig. 4. Convex hull boundary for (a) one connected component; and (b) several connected components. The convex hull boundary is represented by a red dashed line. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article).




Fig. 5. Gap detection. (a) Normals are calculated around the convex hull boundary using the c -scale shape descriptor. (b) Thickness of the segmented object is calculated at each boundary element of the convex hull. Zero thickness indicates the presence of a gap.

2.4.2. Object completion

The object completion is performed in three steps: convex hull boundary extraction, gap detection, and gap completion.

First, the convex hull of the segmentation result is computed and its boundary extracted (Fig. 4). If the segmented object is composed by several connected components, the convex hull contains all of them, as shown in Fig. 4(b). In the following, the completion strategy is illustrated on an example with three gaps to describe a general case. In our application, all the objects will have at least one gap to complete.

For each boundary element in the convex hull boundary, the c -scale shape descriptor (Rueda et al., 2008) is used to define the tangent (chord) at each point in the curve. From the tangents, the normals to the convex hull boundary are calculated at each boundary point in the direction pointing towards the inside of the object (Fig. 5(a)). Then, the binary intersection between each normal and the segmented object (resulting from the feature-based fuzzy connectedness step) is retrieved and the connected object closest to the convex hull boundary element is retained. The width (thickness) is then calculated (Fig. 5(b)) by measuring the length of the connected object previously extracted for each boundary element of the convex hull. The gap(s) is(are) detected by finding the region(s) with zero width, as represented in Fig. 6.

The last step consists of filling in the gap(s) in the segmented object. Two normals on each side of each detected gap are then identified at a fixed distance D (Fig. 7(a)). A polygon is constructed by connecting the two detected normals on each side of the hole at the level of the convex hull boundary from one side, and the segmented object boundary from the other side (Fig. 7(b)). The corrected object is obtained from the binary union between the polygon and the segmented object. Note that a different completion strategy could have been used instead of a polygon. For example, curvature information could be used to complete the gaps after detection using

curves. However, in this case, we tried to follow the same strategy clinicians were using to complete the shapes in our particular application. Algorithm 1 summarises the object completion step.

Algorithm 1 Object completion.

Input: Segmented binary image S resulting from the AFC step.

Output: Completed object.

- 1: Compute the convex hull C of S ;
 - 2: Extract the boundary $\mathcal{B} = \{p_i : i = 1, \dots, q\}$ of C ;
 - 3: **for** each $p_i; i = 1 : q$ **do**
 - 4: Calculate the c -scale value $C_h(p_i)$ and c -scale segment $C(p_i)$;
 - 5: Calculate the normal \mathbf{n}_{p_i} at each p_i as the line perpendicular to $C_h(p_i)$ passing through p_i and pointing towards the inside of the object;
 - 6: Calculate the thickness $\mathcal{T}(p_i)$ at p_i as the connected object resulting from the binary intersection between \mathbf{n}_{p_i} and S and closest to \mathcal{B} .
 - 7: **end for**
 - 8: Output $\mathcal{T}(p_i)$ for all i ;
 - 9: Detect the gap(s) $G_j, j = 1 : m$, as region(s) where $\mathcal{T}(p_i) = 0$;
 - 10: **for** each $G_j; j = 1 : m$ **do**
 - 11: Identify the two normals \mathbf{n}_f and \mathbf{n}_b at both sides of G_j at a fixed distance D from G_j ;
 - 12: Construct a polygon by connecting the lines \mathbf{n}_f and \mathbf{n}_b at the level of C from the outside and the segmented object in S from the inside;
 - 13: **end for**
 - 14: Output completed object as the binary union between the polygon(s) and the segmented object in S .
-

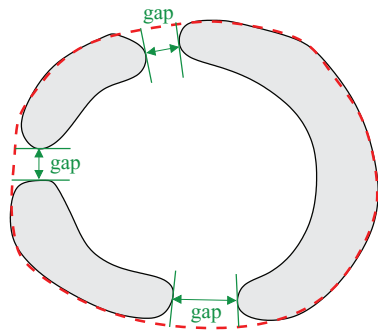


Fig. 6. Gaps detected after identification of zero thickness regions as shown in Fig. 5.

2.5. Regularisation

The resulting object boundary is finally smoothed using a mean curvature flow (MCF)¹ regularisation strategy (Sethian, 1999). The method is based on the evolution of the curve using implicit functions. The points in the contour are moved in the normal direction with a speed proportional to the curvature at each point. A Matlab toolbox (Mitchell, 2008) was used for this purpose.

2.6. Implementation details

Local phase and feature asymmetry were estimated as described in Sections 2.1 and 2.2. The bandpass filter used within this framework is a Gaussian derivative filter, as defined in Eq. (3). Since the scale considered for local phase calculation (Eq. 6) depends on the size of the structure of interest, two different scales were considered, one for gestational ages below 30 weeks ($s = 27$), and one for gestational ages above or equal to 30 weeks ($s = 35$). For the calculation of the feature asymmetry (Eq. 8), three scales were considered ($N = 3$), with $s = [23, 25, 27]$ for gestational ages below 30 weeks and $s = [27, 30, 35]$ for gestational ages above 30 weeks. T_s was obtained from statistical properties of the local phase image (Kovesi, 1999), and set to $T_s = 0.155$.

Within the AFC framework (cf. Section 2.3), for the object feature-based component of affinity (11), the mean $m_o = 2.44$ and the standard deviation $\sigma_o = 3 \times 0.086$ were estimated from a region of fat in the local phase image from a training stage performed on three images, since the local phase value of a region of fat was very similar among images and didn't require a larger training set. The images used for training were not part of the evaluation set. The final affinity (13) was calculated with $\omega_1 = \omega_2 = 0.5$. The method is multi-seeded

with one or more seeds in the fat layer of the image used for initialisation. In the set of images used, most of the images required one seed and a few required more, if presenting disconnected fat appearance. No more than 5 seeds were used in any of the images. The object of interest was thresholded from the connectivity map by using $T_{FC} = 0.85$. This value was set empirically but could have been automated for each seed as in (Miranda et al., 2008).

The c -scale shape descriptor used for delineating closed regions of adipose tissue (cf. Section 2.4) only required one parameter t , which was set to $t = 5$.

The AFC part of the method was implemented in Matlab, using C mex files for faster computation. The other steps of the presented framework were implemented in Matlab.

3. Results and evaluation

This section presents results of evaluation of the new segmentation method on a large clinical dataset. We begin by presenting the clinical image protocol in Section 3.1. The proposed framework is then directly compared to the original Absolute Fuzzy Connectedness method based on intensities and qualitatively and quantitatively against manual segmentations in Sections 3.2 and 3.3, respectively. We then, in Section 3.4, look more deeply at the performance of the algorithm by firstly characterizing the variability of the data in the clinical dataset, and use this characterization to gain better understand into the potential performance of the new segmentation method on real world clinical data.

3.1. Image acquisition

The clinical dataset used for evaluation is 81 cross-sectional US images of the fetal arm across gestation acquired perpendicularly to the arm at mid-humeral level (Fig. 8) from 73 healthy fetuses between 20 and 36 weeks of gestation. For eight of these fetuses, images acquired at two different gestational ages were included. The distribution of gestational ages within the dataset is shown in Fig. 9.

The images were acquired with a Philips HD9 machine (Philips Ultrasound, Bothell, WA, USA) at the Nuffield Department of Obstetrics and Gynaecology, John Radcliffe Hospital, University of Oxford, Oxford, U.K. The fetuses involved in this clinical study are part of the INTERGROWTH-21st (2009)² and INTERBIO-21st (2012)³ cohorts.

The protocol used for the acquisition of the US fetal arm cross-sections was as follows. First, the sagittal view of the humerus (Fig. 10(a)) was acquired to visualise the full humerus length longitudinally, ideally horizontal and in the centre of the screen. The probe

¹ <http://www.cs.ubc.ca/~mitchell/ToolboxLS/>

² <http://www.intergrowth21.org.uk>

³ <http://www.interbio21.org.uk>




Fig. 7. Gap completion using polygons. (a) Normals are found at either side of the gap at a fixed distance D . (b) A polygon is constructed using the normals and the segmented objects. The polygons are represented by green dashed lines. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article).




Fig. 8. US cross-section acquisition of the fetal arm at mid-humeral level (purple cross-section).

was subsequently rotated 90° to obtain an axial cross-section of the arm at mid-humeral level (Fig. 10(b)).

Referring to Fig. 11, arm cross-sections are formed by a central hyperechoic bone surrounded by hypoechoic muscle and then an echodense fat layer. To ensure that the cross-sections were acquired perpendicular to the humerus, the probe was swept along the longitudinal axis of the humerus bone. If the axial view appeared to be perpendicular to the longitudinal axis then the image of the bone remained in the centre of the screen as the probe was moved. Adjustments were made until this was achieved and then returned to midpoint of humerus to acquire the 2D image.

Image appearance was found to vary across gestation as illustrated in Fig. 12. The following general observations can be made to illustrate some of the challenges in image segmentation for this particular application. First, the shape of the fetal arm is not always circular and can vary globally or regionally due to the pressures created by surrounding structures - this is especially the case at later gestational ages. Second, the adipose tissue layer can produce pronounced intensity inhomogeneities, which are characteristic of this imaging modality. Changes in tissue texture can also create different

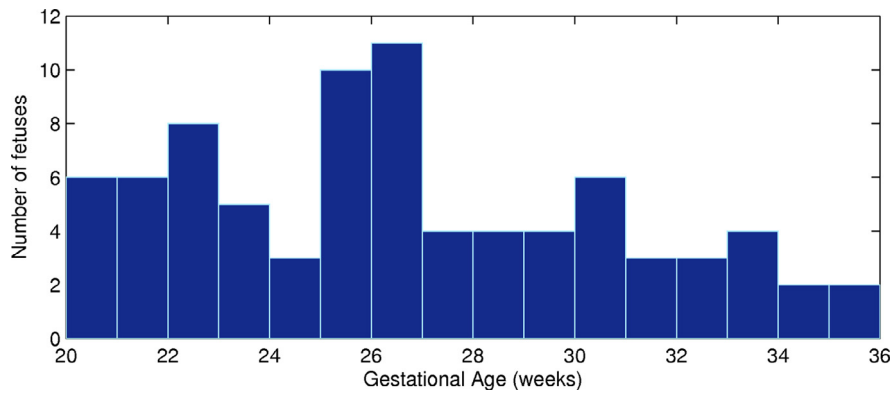


Fig. 9. Distribution of gestational ages within the dataset.




Fig. 10. Fetal arm at 28 weeks of gestation. (a) Sagittal view with horizontal humerus bone. (b) Axial cross-section.




Fig. 11. (a) Schematic of arm composition. (b) Arm cross-section of a 27 weeks fetus with characteristic shadow under the humerus bone, due to lack of ultrasonic signal response from that region. Notice intensity inhomogeneities within the adipose tissue (B).




Fig. 12. Image appearance of the fetal arm US cross-sections across gestational age.

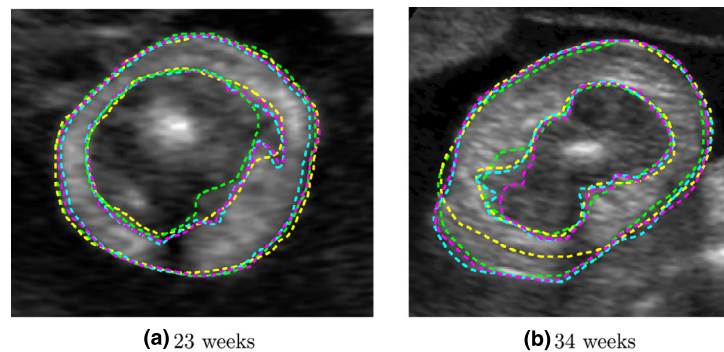


Fig. 13. Manual segmentations performed twice by two different experts. (Expert 1: green and yellow contours. Expert 2: magenta and cyan contours). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article).

speckle patterns at different gestational ages. Third, maternal and fetal tissues (e.g. as shown in Figs. 12(d, f and h)) normally surround the arm, and can make the segmentation task difficult. Fourth, at early gestation, the layer of fat is very thin and hardly visible, since it is seen more clearly from 18–20 weeks. Fifth, the adipose tissue boundaries are usually fuzzy, which makes manual segmentation difficult and can cause discrepancies, as shown in Fig. 13, where the adipose tissue layers were manually segmented by two experts twice. Finally, observe that there is always a characteristic shadow appearing under the humerus bone (Figs. 11(b) and 12), which prevents the visualization of adipose tissue in that area. In manual segmentations, this region is typically approximated by joining the delineations on either side of the shadow by a straight line (Fig. 13).

3.2. Qualitative evaluation

In this subsection, the proposed method is compared to the intensity-based Absolute Fuzzy Connectedness approach. However, due to the large intensity variability within the adipose tissue layer across the different images in our dataset, it proved impossible to set representative algorithm parameters for the intensity-based AFC method during the training stage. This situation is avoided when using local phase, as it is contrast invariant. A typical example of intensity-based segmentation is shown in Fig. 14. We observed that the intensity-based approach could not cope with the inhomogeneities present within the object of interest. This can be seen in Fig. 14(c), where high intensity regions within the adipose tissue area

are not segmented. In this case, the variability of intensities within the region of interest is too high for the intensity-based method to correctly segment the overall adipose tissue layer.

Fig. 15 shows the outputs of key image analysis steps of the proposed methodology reported in this paper. Qualitative results comparing the automated method output with manual delineations at a number of discrete gestational ages are shown in Fig. 16. These results illustrate that the automated method appear visually similar to the manual delineations.

3.3. Quantitative evaluation

In this subsection, we quantitatively assess the proposed segmentation method by using a number of established region-based and distance-based metrics. First, region-based evaluation metrics, defined as area overlap measures, were selected as a way of assessing image segmentation precision (repeatability of the method) and accuracy (sensitivity and specificity). These metrics are as defined in Udupa et al. (2006). Experimental results were performed twice on each image of the dataset to assess the precision of the proposed method. Accuracy was reported as in Udupa et al. (2006), where delineation *sensitivity* is given by the true positive area fraction (TPAF) and delineation *specificity* by $1 - \text{FPAF}$ where FPAF is the false positive area fraction. These two independent metrics are sufficient to quantify the general accuracy of a segmentation method. In each case, a larger value indicates a better segmentation performance.




Fig. 14. Intensity-based FC segmentation results. (a) Arm cross-section of a 28 weeks fetus. (b) Intensity-based FC connectivity map. (c) Segmentation for $T_{FC} = 0.75$. (d) Segmentation for $T_{FC} = 0.9$. Dashed lines: averaged manual segmentation. Continuous lines: FC segmentation results.

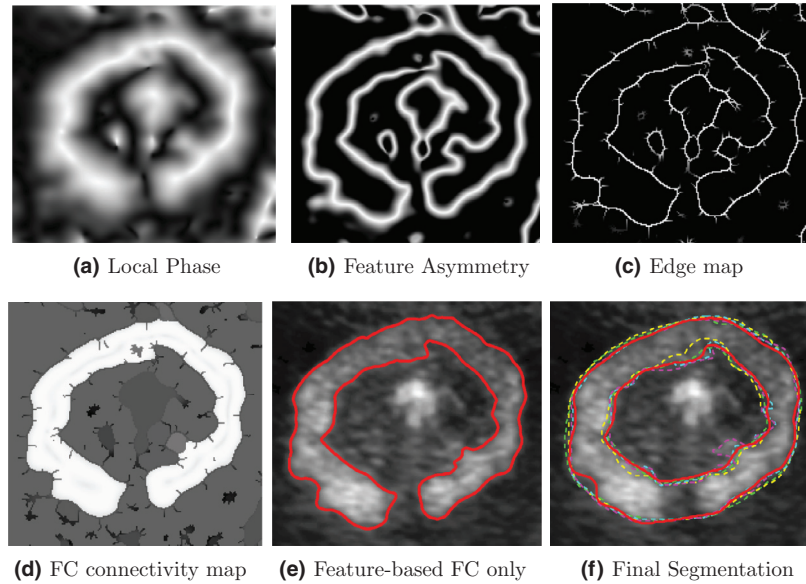


Fig. 15. Main steps of the proposed method applied to Fig. 11(b). (a) Local Phase at $\sigma = 27$. (b) Feature asymmetry with $\sigma = 23$, $\sigma = 25$, $\sigma = 27$. (c) Edge map extracted from feature asymmetry using non-maximal suppression in all directions. (d) FC connectivity map. (e) Feature-based FC segmentation result. (f) Final segmentation obtained from feature-based FC after completion and regularisation. Manual segmentations are displayed in dashed lines. Continuous lines show the proposed segmentation results.

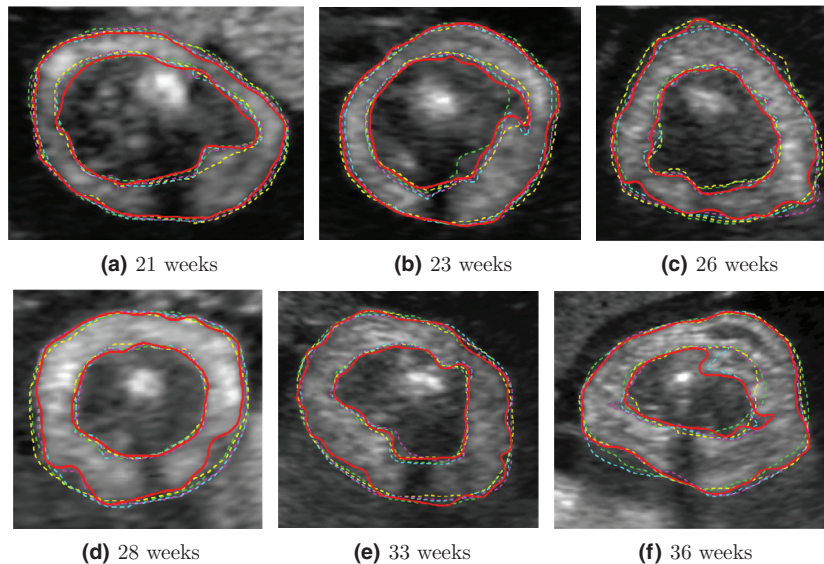


Fig. 16. Segmentation results across gestational ages. Manual segmentations are displayed in dashed lines. (Expert 1: green and yellow lines. Expert 2: magenta and cyan lines.) Continuous red lines show the proposed segmentation results. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article).

We also report the commonly used Dice similarity metric. Distance-based metrics (maximum symmetric contour distance: MSD; average symmetric contour distance: ASD; and root mean square symmetric contour distance: RMSD), as described in Heimann et al. (2009), are also reported. As we do not have a “ground-truth” segmentation (the true arm composition is not known but only im-

aged indirectly), segmentation results were compared to manual delineations of the structures, segmented twice by each of the two experts. The results per image were averaged to obtain the overall performance for a particular expert and for all experts. More details on these particular metrics can be found in Rueda et al. (2014).

Table 1

Intra and inter-observer variability. Manual delineations from Expert 1 (E1) and Expert 2 (E2) are evaluated against themselves and against each other using area overlap and distance-based metrics. The area overlap metrics evaluated are precision, accuracy (sensitivity and specificity), and Dice similarity as defined in Udupa et al. (2006). The distance-based metrics evaluated are the maximum symmetric contour distance (MSD), the average symmetric contour distance (ASD), and the root mean square contour distance (RMSD) as defined in Heimann et al. (2009).

	Intra-expert variability		Inter-expert variability	
	E1	E2	E1 vs E2	
	Precision (%)	83.49 ± 4.10	87.06 ± 3.06	80.29 ± 3.99
Accuracy (%)	Sensitivity	90.15 ± 4.75	94.19 ± 2.72	88.10 ± 5.29
	Specificity	98.11 ± 0.95	98.11 ± 0.93	97.59 ± 1.38
Dice (%)	90.95 ± 2.46	93.05 ± 1.77	88.99 ± 2.49	
MSD (mm)	1.02 ± 0.52	0.93 ± 0.49	1.27 ± 0.68	
ASD (mm)	0.29 ± 0.13	0.23 ± 0.10	0.36 ± 0.16	
RMSD (mm)	0.38 ± 0.18	0.31 ± 0.15	0.47 ± 0.22	

Table 2

Quantitative evaluation. Automatic segmentations (auto) are evaluated against the ground truth, generated from manual delineations from Expert 1 (E1) and Expert 2 (E2), using area overlap and distance metrics. The area overlap metrics evaluated are accuracy (sensitivity and specificity) and Dice similarity as defined in Udupa et al. (2006). The distance-based metrics evaluated are the maximum symmetric contour distance (MSD), the average symmetric contour distance (ASD), and the root mean square contour distance (RMSD) as defined in Heimann et al. (2009).

		Auto vs E1	Auto vs E2	Mean
Accuracy (%)	Sensitivity	85.63 ± 4.55	88.98 ± 4.38	87.30 ± 3.84
	Specificity	96.86 ± 1.41	97.23 ± 1.12	97.05 ± 1.17
Dice (%)		86.02 ± 2.90	88.21 ± 2.79	87.11 ± 2.60
MSD (mm)		1.72 ± 0.86	1.65 ± 0.86	1.68 ± 0.82
ASD (mm)		0.46 ± 0.19	0.36 ± 0.18	0.41 ± 0.18
RMSD (mm)		0.58 ± 0.24	0.49 ± 0.24	0.54 ± 0.23

Table 1 presents the results for the intra and inter-observer variability assessment, obtained from the manually segmented images by two experts, segmented twice. The results show similar performance between experts, Expert 2 having slightly better results.

The segmentation evaluation results of the proposed approach were then compared to both experts and to the average manual segmentation, as shown in Table 2.

The proposed method performs similarly to manual delineation with mean results very close to those obtained for each metric calculated for the inter-expert variability (cf. Table 1) in terms of mean and standard deviation. The precision of the proposed segmentation approach, in terms of repeatability, was evaluated by repeating each segmentation twice using different seed locations as initialisation. The presented framework has a precision of 99.89%, which means that the results are very consistent. Very slight differences were noted in certain cases due to the selection of the seed positions. The repeatability is much higher than the one obtained from manual adipose tissue delineations (cf. Table 1) as expected.

3.4. Quantitative image quality assessment

This subsection firstly explains how we define image segmentation quality for our dataset and then interprets the automated algorithm performance with respect to the resulting image segmentation quality metrics.

It would be greatly beneficial to report segmentation results with a measure of image quality to characterise the dataset used and how well the method performs considering the quality of the images.

However, establishing overall image quality measures is difficult, since the quality of images relies on tissue appearance. Ultrasound image quality can vary considerably between acquisitions, which may affect the performance of different segmentation methods.

In this paper, we propose a new solution to quantify image quality of a clinical dataset designed to provide deeper insight into segmentation performance. This is, to our knowledge, the first attempt to correlate segmentation results with a quantitative measure of US image quality.

US image quality depends on a number of factors including: the US machine (transducer, time-gain control, use of harmonics versus fundamental, persistence, and depth), the object being scanned (tissue properties (speckle), effects of attenuation (depth), shadows, and reverberations), and the orientation of the probe with respect to the object.

In fetal ultrasound imaging, object appearance varies with gestational age with the structures surrounding the object of interest showing high variability. Overall fetal US image quality tends to decrease towards later gestation as a result of the fetus becoming bigger with relatively less amniotic fluid, thus the fetal structures are more likely to be compressed resulting in the clear soft tissue/fluid interface. The bone density in the fetus also increases, creating more shadows and artefacts in the images. Another factor that can affect ultrasound image quality is the increase of maternal body mass index, attenuating the signal especially towards the end of pregnancy. Specifically, the proposed quantitative image quality assessment method relies on the principle that different tissues have specific sound propagation properties characterised by the complexity




Fig. 17. Fetal arm image partitioning into 4 regions (cyan: background region, white: adipose tissue layer, magenta: muscle region, and yellow: bone region) for a 27 week fetus. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article).

of the speckle pattern. These tissues do not evolve in the same way across gestation and surrounding structures vary depending on the acquisition angle and fetal position at that particular time. This is why an overall global image measure would not be as appropriate as a regional quality measure.

Traditional image processing quality measures such as SNR (signal-to-noise ratio) and CNR (contrast-to-noise ratio) rely on the estimation of the signal and noise from the entire image or regions. In the case of ultrasound image processing, and particularly the literature on speckle-reduction, speckle has typically been treated as the noise component. In our case a speckle-based measure would only capture texture differences, not contrast changes, and requires access to the RF signal to estimate a statistical model, for instance see Destremes and Cloutier (2010); Raju and Srinivasan (2002). Further, the estimation of such models is non-trivial with accuracy depending on the block size used for parameter estimation for instance (Larrue and Noble, 2014). In our case, which is very typical of most image analysis work conducted with clinical groups, we have access to DICOM B-mode images only. Furthermore, due to the stochastic nature of the speckle patterns, using CNR directly for characterising echogeneity is sub-optimal, because the tissue contrast resolution depends on speckle variance and size. As these general image quality measures did not satisfy our needs we developed the approach described next.

The proposed new method quantifies the complexity of each region (resulting from the speckle distribution) and the relationship amongst tissues in an image without estimating a speckle model (statistical distribution). Specifically, first a manual image partitioning is made to each image resulting in manual delineations of different regions of interest. An entropy-based measure is then computed on each of the image partitions to estimate the information content in each region of interest. This quality measure is based on the appearance and complexity of each region and not on contrast, absolute

intensity, or edge information. The probability density function of a region, denoted p_r , is first estimated from the gray-level histogram of that region. The normalised histogram of a region A_r is defined for each intensity value a_k with $k = 1, \dots, M$, M being the maximum number of intensity levels in A_r . In our case, $M = 256$. The entropy H of the random variable A_r can then be calculated as

$$H(A_r) = - \sum_{k=1}^M p_r(a_k) \log_2 p_r(a_k). \quad (14)$$

The entropy difference between adjacent regions can then be calculated to assess the overall image quality (as a whole complexity measure) and correlated with the segmentation results.

The proposed quantitative image quality assessment method was applied to the fetal arm dataset introduced in Section 3.1. The first step consisting of partitioning the images into different areas, is shown in Fig. 17. These regions were manually delineated in all the images of the dataset and the entropy calculated for each of these regions separately. The ideal image appearance, from an automated segmentation algorithm perspective, occur when the background and muscle regions have a hypoechoic appearance (dark) and the adipose tissue layer a hyperechoic appearance (bright), which should be clearly distinguishable from the surrounding tissues.

In the evaluation dataset used within this study, we deliberately (and unusually) selected examples with a wide range of image quality. The examples discussed below are typical examples taken from the whole dataset and were chosen to facilitate the understanding of how entropy values relate to the fetal image regions analysed. We have studied in detail the relationship between all the image regions in the dataset and the entropy values before concluding how to generalise our findings which are reported below.

The entropy of the background region is shown in Fig. 18 for all the images in the dataset. A selection of representative images with low, medium, and high background entropy values in Fig. 18 are shown in Fig. 19 to visually appreciate the difference. Observe that higher entropy values are correlated with the presence of more fetal and maternal tissues surrounding the adipose tissue layer. The lower the entropy, the clearer the interface between background and adipose tissue.

Similarly, the entropy for the adipose tissue region (cf. white region in Fig. 17) across gestational age is shown in Fig. 20, with representative examples for low, medium, and high entropy shown in Fig. 21.

The adipose tissue regions in Fig. 21(a and d) present more information, showing higher intensity levels in these regions. Regions outlined in Fig. 21(c and f) have lower entropy, and their appearance looks fuzzier, visually corresponding to a lower quality. Ideally, we would like the adipose tissue region to be associated with high entropy.

The entropy values for the muscle region (cf. magenta region in Fig. 17) are represented in Fig. 22 across gestation. Examples for low, medium, and high entropy, as indicated in Fig. 22, are shown in




Fig. 18. Entropy of the background region across gestational age for all the images in the dataset. The numbers within the coloured bullets correspond to the images in Fig. 19.




Fig. 19. Examples of background entropy as colour-coded in Fig. 18. (a,d) High entropy; (b,e) medium entropy; and (c,f) low entropy. The higher the entropy, the more fetal and maternal tissues surround the adipose tissue layer.




Fig. 20. Entropy of the adipose tissue region across gestational age. The numbers within the coloured bullets correspond to the images in Fig. 21. High entropy denotes better image appearance for the adipose tissue region.

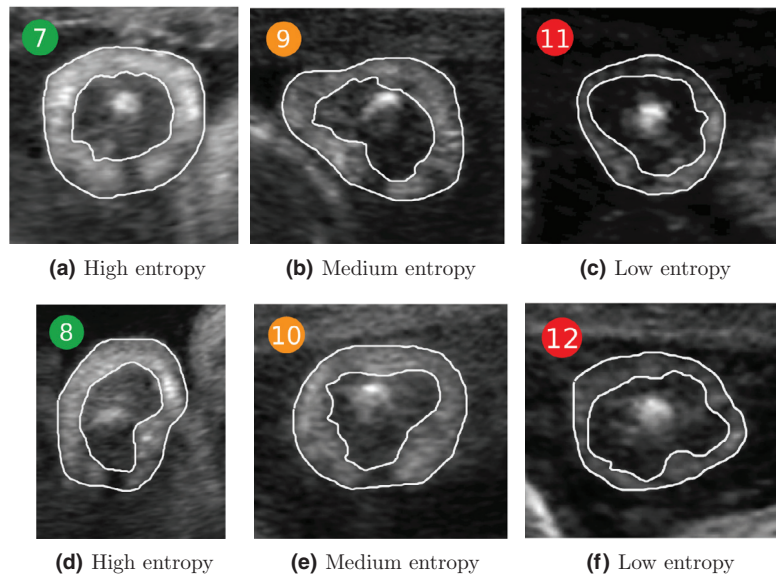


Fig. 21. Examples of adipose tissue entropy as colour-coded in Fig. 20. (a,d) High entropy; (b,e) medium entropy; and (c,f) low entropy.




Fig. 22. Entropy of the muscle region across gestational age. The numbers within the coloured bullets correspond to the images in Fig. 23. Ideally, we would like the muscle region to have low entropy and a dark appearance.

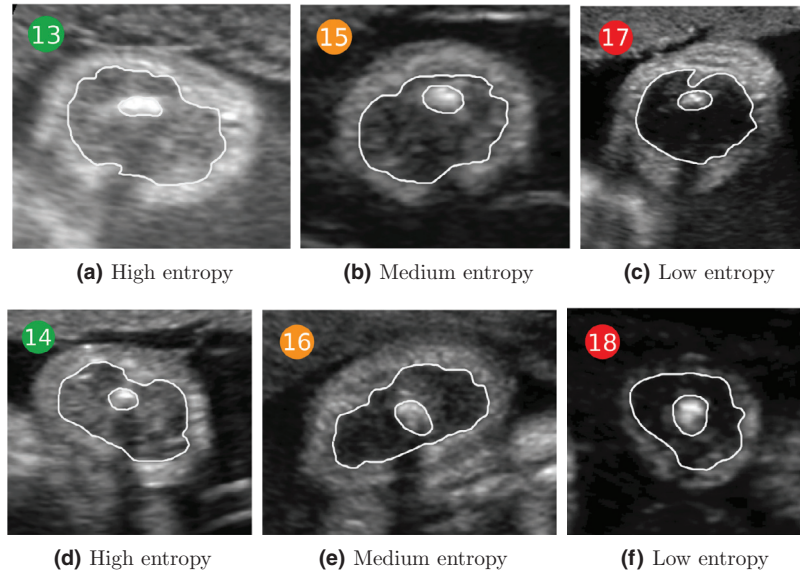


Fig. 23. Examples of muscle entropy as colour-coded in Fig. 22. (a, d) High entropy; (b, e) medium entropy; and (c, f) low entropy.

Fig. 23. Observe that low entropy muscle regions (Fig. 23(c and f)) are much darker than high entropy muscle regions (Fig. 23(a and d)), showing that the information content in these regions is very different. Ideally, we would like the muscle region to be dark, hence to have low entropy.

The entropy for the humerus bone region (cf. yellow region in Fig. 17) is represented in Fig. 24 across gestation. Representative examples for low, medium, and high entropy, as indicated in Fig. 24, are shown in Fig. 25. In this case, the difference is not as noticeable as for the other regions, due to the small size of the structure. However, it can be seen that regions with lower intensity variability within the region have lower entropy, as shown in Fig. 25(c and f).

Comparing the 4 regions, the highest mean entropy is observed for the humerus bone region, with a value of 6.40 bits. Then, the adipose tissue mean entropy has a value of 6.11 bits, the background region has a mean entropy value of 5.38 bits, and the muscle region has a mean entropy value of 5.27 bits. We conclude from this that, on average, the bone presents the highest information content, followed by the adipose tissue region. Background and muscle regions have lower information content.

Having looked at the entropy (and entropy variation across gestational age) for different tissues of interest we now consider how to define image segmentation quality metrics. Recall, that the goal is to segment the adipose tissue layer. Thus the two interfaces of interest are background–adipose tissue, and adipose tissue–muscle. Therefore, we define two scores to assess the difference of entropy between the background region and the adipose tissue region, and the adipose tissue region and the muscle region for each

image in the evaluation dataset. Let S_{ab} be the score representing the difference in entropy between adipose tissue and background, defined as

$$S_{ab} = H(A_{\text{adipose tissue}}) - H(A_{\text{background}}); \quad (15)$$

and S_{am} the score associated with the difference of entropies between adipose tissue and muscle, defined as

$$S_{am} = H(A_{\text{adipose tissue}}) - H(A_{\text{muscle}}), \quad (16)$$

with $H(A_{\text{adipose tissue}})$, $H(A_{\text{background}})$, and $H(A_{\text{muscle}})$ as defined in (14). Both scores are useful in assessing the overall image segmentation quality, as shown in Fig. 26, where each value is colour-coded by its corresponding gestational age.

Fig. 26 represents the image quality of each image by using the scores S_{ab} and S_{am} , derived from adjacent regions as previously defined. Score values are low when adjacent regions are similar, and hence image quality overall is lower. Higher score values translate into more distinct adjacent regions, and hence higher image quality overall. Gestational age is incorporated into Fig. 26, as it is normally correlated to image quality (the image quality generally decreases with gestational age). This dataset was chosen to be representative of this particular application, showcasing a variety of image qualities as found in clinical practice. The entropy-based analysis shows that the dataset has a correspondingly high variability in terms of entropy, including several cases with negative S_{ab} values, where the entropy in the background is higher than the information of the adipose tissue layer. This can happen when the arm is surrounded by other organs (e.g. limbs, abdomen, placenta) as shown in Fig. 19 for




Fig. 24. Entropy of the humerus bone region across gestational age. The numbers within the coloured bullets correspond to the images in Fig. 25.




Fig. 25. Examples of bone entropy as colour-coded in Fig. 24. (a, d) High entropy; (b, e) medium entropy; and (c, f) low entropy.




Fig. 26. Image quality assessment for the fetal arm dataset. S_{ab} denotes the difference of entropy between adipose tissue and background, whereas S_{am} represents the difference of entropy between adipose tissue and muscle. Each value has been colour coded with its corresponding gestational age given in weeks. The lower the score values, the more similarity between adjacent regions, and the lower the image quality (bottom left hand side corner of the graph). The higher the score values, the more difference between adjacent regions, and the higher the quality (top right hand side corner of the graph).




Fig. 27. Segmentation precision (repeatability) with respect to image quality scores S_{ab} and S_{am} . Varying the initial seeds result in mostly the same segmented object except for a few cases where small differences appear.

the high entropy examples. These effects are important to consider in the assessment of a segmentation method, as they present challenging conditions for any method.

Having explained how to define image segmentation quality metrics for our clinical dataset we can now look at how the automated segmentation method performs with respect to these metrics. Figs. 27 and 28 show how precision and accuracy correlate with the two image segmentation metrics, respectively.

The repeatability of the method is presented in Fig. 27. Observe that in most cases the results obtained are very similar when varying the position of the initialisation seeds. Only small differences can be observed in a few cases (values below 1) across the whole range of image qualities.

As shown in Fig. 28 the proposed segmentation method performs robustly over a range of image qualities, giving high values of accuracy in most cases, independently of their image appearance. The lowest accuracy values (blue colours) in terms of sensitivity occurred for images in the bottom left hand side quadrant of Fig. 28.(a), where more similarity between adjacent regions exists (lower quality since adjacent tissue layers look similar) and where the background seems to present more surrounding structures (negative S_{ab} values). However, high precision and accuracy values (red and yellow colours) can also be observed in that same quadrant. It is worth pointing out that the lowest accuracy values observed are above 80%, which is good in terms of segmentation performance. Therefore, we conclude that the proposed segmentation method is robust across the variety of image qualities present in the clinical evaluation dataset.

4. Discussion and conclusions

This paper has presented three main technical contributions: a feature-based segmentation strategy adapted to US images, a gap completion method, and a novel quantitative image quality assessment approach to assess segmentation performance.

The complete US image segmentation framework introduced in this paper is based on a feature-based fuzzy connectedness segmentation method and requires manual placement of the seeds, after which the remaining steps are performed automatically. The selection of the threshold was fixed for this application, in future might be automated by, for instance, the method of Miranda et al. (2008). The proposed approach uses structural and edge information based on local phase, instead of intensities and intensity gradients, to drive

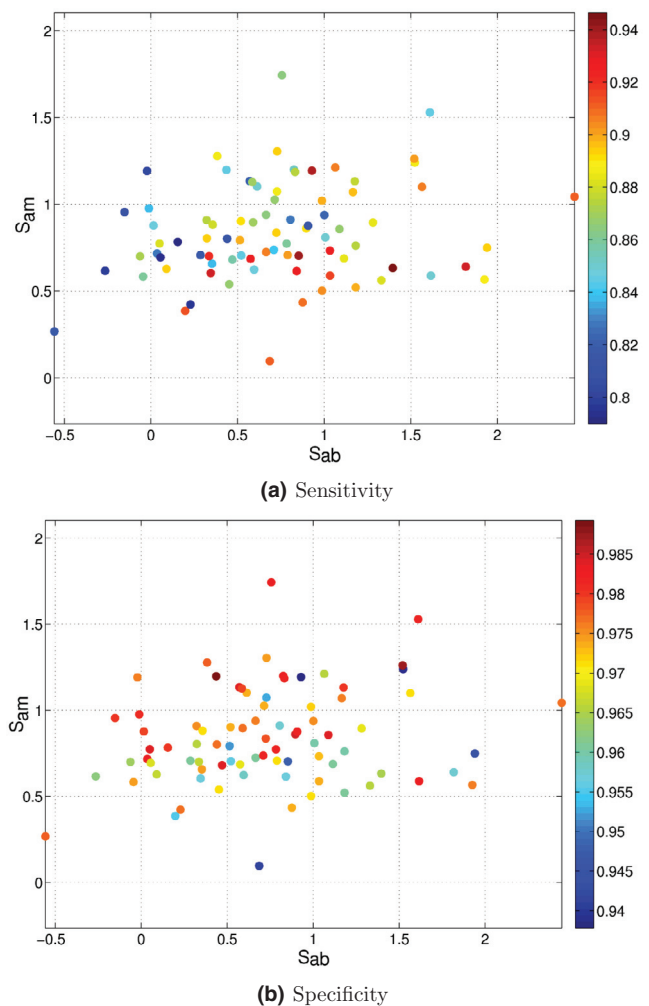


Fig. 28. Segmentation accuracy with respect to image quality scores S_{ab} and S_{am} . (a) Sensitivity. (b) Specificity. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article).

the segmentation. The resulting segmentation is then completed by filling one or more gaps caused by shadows or artefacts in the segmented object of interest using a shape descriptor. A final regularisation based on mean curvature flow is performed to smooth the final contours.

Although more conceptually advanced fuzzy connectedness methods exist, such as RFC and IRFC, it remains to be seen how these would perform in US images. This paper reports results on AFC applied to US images, which is the most basic form of FC with affinities specially formulated for US image segmentation. Once this basic investigation is reported and the behaviour of AFC understood in its most fundamental form, we can then take on investigations to study how more advanced forms of FC with the same forms of affinities on a multi-object setting would perform.

We argued that all segmentation methods should report their results in conjunction with a quantitative image quality analysis to show that the dataset used is representative of a clinical application, and not selected to best suit a particular methodology. A novel quantitative image quality assessment protocol based on entropy was presented and applied to different image partitions to derive interface scores to show the variability of qualities existing in the dataset, representative of a real clinical application. This technique could readily be adapted to suit images from different clinical applications.

A qualitative and quantitative evaluation was performed on 81 cross-sectional images of the fetal arm across gestation, by using region and distance-based metrics. The results showed a similar performance to manual segmentations. Furthermore, the quantitative image quality assessment method showed that the performance of the method was robust across a variety of image qualities representative of a real clinical environment.

The proposed method has undergone clinical assessment on pilot data (Knight et al., 2012a; 2012b; Rueda et al., 2012a) and is now part of a large clinical study aimed at establishing normative nutritional growth charts of healthy fetuses across gestation (Knight et al., 2014). The presented framework estimates three main clinical measurements from US images: the amount of fetal arm adipose tissue, the fat-free (lean and bone) areas (useful for body composition assessment), and the adipose tissue percentages for each cross-section (normalised measurements with respect to arm size) across gestational ages. In this study, we have analysed cross-sectional data, but the method is also suitable to study longitudinal data, towards achieving a personalised nutritional monitoring of the fetus.

The 2D feature-based FC implementation could readily be extended to 3D, as local phase and fuzzy connectedness can be easily extended to 3D. Finally, the proposed framework is motivated by, but not limited to this particular application or imaging modality and could equally be applied to other soft tissue segmentation problems, such as myocardium segmentation (Dietenbeck et al., 2012; Zhu et al., 2010), including contrast-enhanced US (CEUS) images, or intravascular US (IVUS) (Ciompi et al., 2012; Moraes and Furuie, 2011; Zhu et al., 2011).

Acknowledgements

This work was funded as part of the Oxford Centre of Excellence in Personalised Healthcare, funded by the Wellcome Trust and EPSRC Centres of Excellence in Medical Engineering scheme (grant WT 088877/Z/09/Z). AP also acknowledges the Oxford Partnership Comprehensive Biomedical Research Centre funded by the Department of Health NIHR Biomedical Research Centres funding scheme. Clinical data acquisition was approved under the Intergrowth-21st project ethics approval (Oxfordshire Research Ethics Committee C ref 08/H0606/139). Special thanks to Prof. J.K. Udupa and Dr. A.R. Cifor for their useful comments and suggestions.

References

- Belaid, A., Boukerroui, D., Maingourd, Y., Lerallut, J., 2011. Phase-based level set segmentation of ultrasound images. *IEEE Trans. Inf. Technol. Biomed.* 15 (1), 138–147.
- Bernstein, I., Goran, M., Amini, S., Catalano, P., 1997. Differential growth of fetal tissues during the second half of pregnancy. *Am. J. Obstet Gynecol.* 176, 2832.
- Boukerroui, D., Noble, J., Brady, M., 2004. On the choice of band-pass quadrature filters. *J. Math. Imaging Vis.* 21, 53–80.
- Boukerroui, D., Noble, J.A., Robini, M.C., Brady, M., 2001. Enhancement of contrast regions in suboptimal ultrasound images with application to echocardiography. *Ultrasound Med. Biol.* 27 (12), 1583–1594.
- Bridge, C., Noble, J., 2015. Object localisation in fetal ultrasound images using invariant features. *Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI) 2015: from Nano to Macro*.
- Carneiro, G., Amat, F., Georgescu, B., Good, S., Comaniciu, D., 2008a. Semantic-based indexing of fetal anatomies from 3-d ultrasound data using global/semi-local context and sequential sampling. *IEEE Comput. Vis. Pattern Recognit.* 1–8.
- Carneiro, G., Georgescu, B., Good, S., 2008b. Knowledge-based automated fetal biometrics using syngo Auto OB measurements. *Siemens Med. Solut* 1–5.
- Carneiro, G., Georgescu, B., Good, S., Comaniciu, D., 2008c. Detection and measurement of fetal anatomies from ultrasound images using a constrained probabilistic boosting tree. *IEEE Trans. Med. Imaging* 27 (9), 1342–1355.
- Chalana, V., Winter, T.C., Cyr, D.R., Haynor, D.R., Kim, Y., 1996. Automatic fetal head measurements from sonographic images. *Acad. Radiol.* 3 (8), 628–635.
- Ciesielski, K., Udupa, J., 2010. Affinity functions in fuzzy connectedness based image segmentation ii: defining and recognizing truly novel affinities. *CVIU* 114, 155–166.
- Ciesielski, K., Udupa, J., 2012. A framework for comparing different image segmentation methods and its use in studying equivalences between level set and fuzzy connectedness frameworks. *CVIU* 115, 721–734.
- Ciesielski, K., Udupa, J., ao, A.F., Miranda, P., 2012. Fuzzy connectedness image segmentation in graph cut formulation: A linear-time algorithm and a comparative analysis. *J. Math. Imaging Vis.* 44, 375–398.
- Ciompi, F., Pujol, O., Gatta, C., Alberti, M., Balocco, S., Carrillo, X., Mauri-Ferre, J., Radeva, P., 2012. Holimab: a holistic approach for media-adventitia border detection in intravascular ultrasound. *Med. Image Anal.* 16 (6), 1085–1100.
- Ciurte, A., Rueda, S., Bresson, X., Nedevschi, S., Papageorghiou, A., Noble, J., Bach-Cuadra, M., 2012. Ultrasound image segmentation of the fetal abdomen: a semi-supervised patch-based approach. *Proceedings of the International MICCAI Workshop on Perinatal and Paediatric Imaging: PaPI 2012* 1–8.
- Cuingnet, R., Somphone, O., Mory, B., Prevost, R., Yaqub, M., Napolitano, R., Papageorghiou, A., Roundhill, D., Noble, J., Ardon, R., 2013. Where is my baby? a fast fetal head auto-alignment in 3d-ultrasound. *Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI) 2013: from Nano to Macro* 768–771.
- Deng, Y., Wang, Y., Shen, Y., Chen, P., 2012. Active cardiac model and its application on structure detection from early fetal ultrasound sequences. *Comput. Med. Imaging Gr.* 36 (3), 239–247.
- Destremes, F., Cloutier, G., 2010. A critical review and uniformized representation of statistical distributions modeling the ultrasound echo envelope. *Ultrasound Med. Biol.* 36 (7), 1037–1051.
- Dietenbeck, T., Alessandrini, M., Barbosa, D., D'hooge, J., Friboulet, D., Bernard, O., 2012. Detection of the whole myocardium in 2d-echocardiography for multiple orientations using a geometrically constrained level-set. *Med. Image Anal.* 16 (2), 386–401.
- Dindoyal, I., Lambrou, T., Deng, J., Ruff, C., Linney, A., Rodeck, C., Todd-Pokropek, A., 2005. Level set segmentation of the fetal heart. *Funct. Imaging Model. Heart* 3504, 123–132.
- Felsberg, M., Sommer, G., 2001. The monogenic signal. *IEEE Trans. Signal Process.* 49 (12), 3136–3144.
- Feng, S., Zhou, S.K., Good, S., Comaniciu, D., 2009. Automatic fetal face detection from ultrasound volumes via learning 3d and 2d information. *IEEE Comput. Vis. Pattern Recognit.* 2488–2495.
- Gutiérrez Becker, B., Cosío, F.A., Huerta, M.E.G., Benavides-Serralde, J., 2010. Automatic segmentation of the cerebellum of fetuses on 3d ultrasound images, using a 3d point distribution model. *32nd Annual International Conference of the IEEE EMBS* 4731–4734.
- Hachihaliloglu, I., Abugharbieh, R., Hodgson, A., Rohling, R., 2008. Bone segmentation and fracture detection in ultrasound using 3d local phase features. *Med. Image Comput. Assist. Interv.* 11 (Pt 1), 287–295.
- Hanna, C.W., Youssef, A.B.M., 1997. Automated measurements in obstetric ultrasound images. *ICIP* 3, 504–507.
- Heimann, T., van Ginneken, B., Styner, M.A., Arzhaeva, Y., Aurich, V., Bauer, C., Beck, A., Becker, C., Beichel, R., Bekes, G., Bello, F., Binnig, G., Bischof, H., Bornik, A., Cashman, P.M.M., Chi, Y., Cordova, A., Dawant, B.M., Fridrich, M., Furst, J.D., Furukawa, D., Grenacher, L., Hornegger, J., Kainmiller, D., Kitney, R.L., Kobatake, H., Lamecker, H., Lange, T., Lee, J., Lennon, B., Li, R., Li, S., Meinzer, H.-P., Nemeth, G., Raicu, D.S., Rau, A.-M., van Rikxoort, E.M., Rousson, M., Rusko, L., Saddi, K.A., Schmidt, G., Seghers, D., Shimizu, A., Slagmolen, P., Sorantin, E., Soza, G., Susomboon, R., Waite, J.M., Wimmer, A., Wolf, I., 2009. Comparison and evaluation of methods for liver segmentation from CT datasets. *IEEE Trans. Med. Imaging* 28 (8), 1251–1265.
- INTERBIO-21st, 2012. The interbio-21st study: the functional classification of abnormal fetal and neonatal growth phenotypes. www.interbio21.org.uk.
- INTERGROWTH-21st, 2009. The international fetal and newborn growth standards for the 21st century (intergrowth-21st) study protocol. www.intergrowth21.org.uk.

- Knight, C., Ahmed, M., Edwards, K., Donadono, V., Parry, G., Rueda, S., Noble, J., Papageorghiou, A., 2014. Fetal limb fat and lean volume reference ranges from an optimally healthy population. *Proceedings of the 24th World Congress on Ultrasound in Obstetrics and Gynecology (ISUOG) 44 (S1)*, 259.
- Knight, C., Rueda, S., Noble, J., Papageorghiou, A., 2012a. Fetal arm fat: an in utero marker of body composition? *Proceedings of the 22nd World Congress on Ultrasound in Obstetrics and Gynecology (ISUOG) 40 (S1)*, 106.
- Knight, C., Rueda, S., Noble, J., Papageorghiou, A., 2012b. Fetal arm fat: development throughout pregnancy in selectively screened optimally healthy patients. *Proceedings of the 22nd World Congress on Ultrasound in Obstetrics and Gynecology (ISUOG) 40 (S1)*, 105–106.
- Kovets, P., 1999. Image features from phase congruency. *Videre 1 (3)*, 1–26.
- Larciprete, G., Valensise, H., Vasapollo, B., Novelli, G., Parretti, E., Altomare, F., DiPierro, G., Menghini, S., Barbati, G., Mello, G., Arduini, D., 2003. Fetal subcutaneous tissue thickness (sctt) in healthy and gestational diabetic pregnancies. *Ultrasound Obstet. Gynecol.* 22, 591–597.
- Larue, A., Noble, J.A., 2014. Modeling of errors in nakagami imaging: Illustration on breast mass characterisation. *Ultrasound Med. Biol.* 40 (5), 917–930.
- Lei, T., Udupa, J.K., Saha, P.K., Odhner, D., 2001. Artery-vein separation via mra—an image processing approach. *IEEE Trans. Med. Imaging* 20 (8), 689–703.
- Lu, W., Tan, J., 2000. Segmentation of ultrasound fetal images. *Biol. Qual. Precis. Agric.* II 4203 (1), 81–90.
- Lu, W., Tan, J., Floyd, R., 2005. Automated fetal head detection and measurement in ultrasound images by iterative randomized Hough transform. *Ultrasound Med. Biol.* 31 (7), 929–936.
- Mellor, M., Brady, M., 2005. Phase mutual information as a similarity measure for registration. *Med. Image Anal.* 9 (4), 330–343.
- Miranda, P., ao, A.F., Rocha, A., Bergo, F., 2008. Object delineation by k-connected components. *Eurasip J. Adv. Signal Process. (JASP)* 2008, 1–14.
- Mitchell, I., 2008. The flexible, extensible and efficient toolbox of level set methods. *J. Sci. Comput.* 35 (2–3), 300–329.
- Moonis, G., Liu, J., Udupa, J., Hackney, D., 2002. Estimation of tumor volume with fuzzy-connectedness segmentation of mr images. *AJNR* 23 (3), 356–363.
- Moraes, M.C., Furuie, S.S., 2011. Automatic coronary wall segmentation in intravascular ultrasound images using binary morphological reconstruction. *Ultrasound Med. Biol.* 37 (9), 1486–1499.
- Mulet-Parada, M., Noble, J.A., 2000. 2d+t acoustic boundary detection in echocardiography. *Med. Image Anal.* 4 (1), 21–30.
- Namburete, A.L., Noble, J.A., 2013. Fetal cranial segmentation in 2d ultrasound images using shape properties on pixel clusters. *Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI) 2013: from Nano to Macro 2013*, 720–723.
- Namburete, A.L., Ramatullah, B., Noble, J.A., 2012. Nakagami-based adaboost learning framework for detection of anatomical landmarks in 2d fetal neurosonograms. *Ann. Br. Mach. Vis. Assoc. (BMVA)* 2012 (1), 1–16.
- Namburete, A.L., Stebbing, R., Kemp, B., Yaqub, M., Papageorghiou, A., Noble, J.A., 2015. Learning-based prediction of gestational age from ultrasound images of the fetal brain. *Med. Image Anal.* 21 (1), 72–86.
- Namburete, A.L., Yaqub, M., Kemp, B., Papageorghiou, A., Noble, J.A., 2014. Predicting fetal neurodevelopmental age from ultrasound images. *Proceedings of the Medical Image Computing and Computer-Assisted Intervention – MICCAI 2014 LNCS* 8674.
- Ni, D., Yang, X., Chen, X., Chin, C.-T., Chen, S., Heng, P.A., Li, S., Qin, J., Wang, T., 2014. Standard plane localization in ultrasound by radial component model and selective search. *Ultrasound Med. Biol.* 1–15.
- Nithya, J., Madheswaran, M., 2009. Detection of intrauterine growth retardation using fetal abdominal circumference. *ICCTD 2*, 371–375.
- Noble, J.A., Boukerroui, D., 2006. Ultrasound image segmentation: a survey. *IEEE Trans. Med. Imaging* 25 (8), 987–1010.
- Pathak, S.D., Chalana, V., Kim, Y., 1997a. Interactive automatic fetal head measurements from ultrasound images using multimedia computer technology. *Ultrasound Med. Biol.* 23 (5), 665–673.
- Pathak, S.D., Chalana, V., Kim, Y., 1997b. Multimedia systems in ultrasound image boundary detection and measurements. *Proceedings of SPIE Medical Imaging 1997: Image Display* 3031, 397–408.
- Prakash, K.N.B., Ramakrishnan, A.G., Suresh, S., Chow, T.W.P., 2002. Fetal lung maturity analysis using ultrasound image features. *IEEE Trans. Inf. Technol. Biomed.* 6 (1), 38–45.
- Rahmatullah, B., Besar, R., 2009. Analysis of semi-automated method for femur length measurement from foetal ultrasound. *J. Med. Eng. Technol.* 33 (6), 417–425.
- Rahmatullah, B., Papageorghiou, A., Noble, J., 2012. Integration of local and global features for anatomical object detection in ultrasound. *MICCAI 7512*, 402–409.
- Raju, B.I., Srinivasan, M.A., 2002. Statistics of envelope of high-frequency ultrasonic backscatter from human skin in vivo. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* 49 (7), 871–882.
- Rueda, S., Fathima, S., Knight, C.L., Yaqub, M., Papageorghiou, A.T., Rahmatullah, B., Foi, A., Maggioni, M., Pepe, A., Tohka, J., Stebbing, R.V., McManigle, J.E., Ciurte, A., Bresson, X., Bach-Cuadra, M., Sun, C., Ponomarev, G.V., Gelfand, M.S., Kazanov, M.D., Wang, C.-W., Chen, H.-C., Peng, C.-W., Hung, C.-M., Noble, J.A., 2014. Evaluation and comparison of current fetal ultrasound image segmentation methods for biometric measurements: A grand challenge. *IEEE Trans. Med. Imaging* 33 (4), 797–813.
- Rueda, S., Knight, C., Papageorghiou, A., Noble, J., 2011. Local phase-based fuzzy connectedness segmentation of ultrasound images. *Proc. MIUA* 331–335.
- Rueda, S., Knight, C., Papageorghiou, A., Noble, J., 2012a. Novel automatic method for measuring fetal arm adipose tissue in 2d ultrasound images across gestation. *Proceedings of the 22nd World Congress on Ultrasound in Obstetrics and Gynecology (ISUOG) 40 (S1)*, 281.
- Rueda, S., Knight, C., Papageorghiou, A., Noble, J., 2012b. Regularised feature-based fuzzy connectedness segmentation of ultrasound images for fetal soft tissue quantification across gestation. *Proceedings of the 9th IEEE International Symposium on Biomedical Imaging (ISBI) 1323–1326*.
- Rueda, S., Udupa, J., Bai, L., 2008. Local curvature scale: a new concept of shape description. *Proceedings of SPIE Medical Imaging: Image Processing 6914.69144Q1–69144Q11*
- Sethian, J., 1999. *Level Set Methods and Fast Marching Methods*.
- Shrimali, V., Anand, R.S., Kumar, V., 2009. Improved segmentation of ultrasound images for fetal biometry, using morphological operators. *IEEE Eng. Med. Biol. Soc.* 2009, 459–462.
- Sonka, M., Hlavac, V., Boyle, R., 2008. *Image Processing, Analysis, and Machine Vision*.
- Szilágyi, T., Brady, M., Brunner, T., Joshi, N., 2009. Local phase significance estimated with uncertainties to detect fibrotic regions from in vivo pancreatic cancer images. *Proceedings of the 13th Conference in Medical Image Understanding and Analysis, BMVA 204–208*.
- Thomas, J.G., Jeanty, P., Peters, R.A., Parrish, E.A., 1991a. Automatic measurements of fetal long bones. a feasibility study. *J. Ultrasound Med.* 10 (7), 381–385.
- Thomas, J.G., Peters, R.A., Jeanty, P., 1991b. Automatic segmentation of ultrasound images using morphological operators. *IEEE Trans. Med. Imaging* 10 (2), 180–186.
- Udupa, J., Saha, P.K., 2003. Fuzzy connectedness and image segmentation. *Proc. IEEE* 91 (10), 1649–1669.
- Udupa, J., Samarasekera, S., 1996. Fuzzy connectedness and object definition: Theory, algorithms, and applications in image segmentation. *CVGIP* 58 (3), 246–261.
- Udupa, J.K., LeBlanc, V., Zhuge, Y., Imielinska, C., Schmidt, H., Currie, L., Hirsch, B., Woodburn, J., 2006. A framework for evaluating image segmentation algorithms. *Comput. Med. Imag. Grap.* 30 (2), 75–87.
- Udupa, J.K., Nyul, L.G., Ge, Y., Grossman, R.I., 2001. Multiprotocol mr image segmentation in multiple sclerosis: experience with over 1000 studies. *Acad. Radiol.* 8 (11), 1116–1126.
- Veronese, E., Cosmi, E., Visentin, S., Poletti, E., Grisan, E., 2012. Estimation of fetal aorta intima-media thickness from ultrasound examination. *Proceedings of the International MICCAI Workshop on Perinatal and Paediatric Imaging: PaPI 2012* 81–88.
- Yaqub, M., Cuingnet, R., Napolitano, R., Roundhill, D., Papageorghiou, A., Ardon, R., Noble, J., 2013. Volumetric segmentation of key fetal brain structures in 3d ultrasound. *Mach. Learn. Med. Imaging LNCS* 8184, 25–32.
- Yaqub, M., Javaid, M., Cooper, C., Noble, J., 2014a. Investigation of the role of feature selection and weighted voting in random forests for 3d volumetric segmentation. *IEEE Trans. Med. Imaging* 33 (2), 258–270.
- Yaqub, M., Kopuri, A., Rueda, S., Sullivan, P., McCormick, K., Noble, J., 2014b. A constrained regression forests solution to 3d fetal ultrasound plane localization for longitudinal analysis of brain growth and maturation. *Mach. Learn. Med. Imaging LNCS* 8679, 109–116.
- Yu, J., Wang, Y., Chen, P., 2008a. Fetal ultrasound image segmentation system and its use in fetal weight estimation. *Med. Biol. Eng. Comput.* 46 (12), 1227–1237.
- Yu, J., Wang, Y., Chen, P., Shen, Y., 2008b. Fetal abdominal contour extraction and measurement in ultrasound images. *Ultrasound Med. Biol.* 34 (2), 169–182.
- Zhu, X., Zhang, P., Shao, J., Cheng, Y., Zhang, Y., Bai, J., 2011. A snake-based method for segmentation of intravascular ultrasound images and its in vivo validation. *Ultrasonics* 51 (2), 181–189.
- Zhu, Y., Papademetris, X., Sinusas, A., Duncan, J., 2010. A coupled deformable model for tracking myocardial borders from real-time echocardiography using an incompressibility constraint. *Med. Image Anal.* 14 (3), 429–448.