

Multi-Scale Local Phase Features for Anatomical Object Detection in Fetal Ultrasound Images

Bahbibi Rahmatullah¹
bahbibi.rahmatullah@eng.ox.ac.uk

Aris Papageorghiou²
aris.papageorghiou@obs-gyn.ox.ac.uk

J. Alison Noble¹
alison.noble@eng.ox.ac.uk

¹ Institute of Biomedical Engineering
Department of Engineering Science
University of Oxford, UK

² Nuffield Department of Obstetrics &
Gynaecology, John Radcliffe Hospital,
University of Oxford, UK

Abstract

The presence and absence of anatomical landmarks in fetal ultrasound image are commonly used as indicators in finding the optimal ultrasound plane for fetal biometric measurement. We propose a machine learning framework for the automated detection of two important anatomical landmarks (stomach and umbilical vein) in fetal abdominal ultrasound images. Our proposed approach combines the unary features extracted from the multi-scale local phase image with the Haar features acquired from the intensity image. The results presented indicate an improvement over prior intensity-based approach.

1 Introduction

Comprehensive ultrasound examination during pregnancy is required for the clinical purpose of dating pregnancy, monitoring the growth of the fetus and detecting any abnormalities. During the scan, the sonographer scans the thigh, abdomen and head area to get the standard fetal biometric measurements (e.g. the biparietal diameter (BPD), occipito-frontal diameter (OFD), head circumference (HC), abdominal circumference (AC), and femur length (FL)). These biometric measurements are then plotted against the pre-determined growth chart in order to check the normality in the fetal growth. It is very important to have accurate measurements in order to avoid the risks of a false negative (failure to detect abnormality in growth) or false positive which might lead to medical intervention and unnecessary maternal anxiety.

The accuracy of the measurements depend on the correct positioning of the on-screen measurement calliper in the image acquired from a standard plane. The optimality of the acquired image can be checked by identifying the visibility of certain anatomical structures inside the fetus. For example, in a standard fetal abdominal ultrasound scan, the stomach bubble (SB) is visible and the umbilical vein (UV) is at the one-third position from the abdomen wall [1].

The work presented in this paper investigates the effect of introducing a new feature sets into the machine learning framework for the localization of the two important anatomical landmarks (stomach and umbilical vein) in fetal abdominal ultrasound images.

Previous work considered only features from intensity images [2]. Our approach uses the combination of features extracted from multi-scale local phase image and Haar features acquired from intensity images in the training of the object detector using an adaptive boosting (AdaBoost) learning algorithm.

2 Local Phase Features

In ultrasound image analysis, the feature information extracted from local phase-based processing of images has proven beneficial for a variety of image analysis task [3, 4]. Local phase is invariant to changes in image brightness or contrast, making it particularly suitable for ultrasound images, known for its relatively low signal-to-noise ratio, artefacts and shadowing. These works show that a local-phase based method outperforms the conventional intensity-based methods for feature detection in ultrasound images.

The local phase (φ) is defined as:

$$\varphi(x) = \arg(f_A(x)) = \arctan\left(\frac{\text{Im}(f_A(x))}{\text{Re}(f_A(x))}\right) \quad (2)$$

where $f_A(x)$ is the complex analytic signal, $\text{Re}[\cdot]$ and $\text{Im}[\cdot]$ correspond to the real and complex portion of the signal, respectively.

In this implementation, we used an approach known as the monogenic signal [5]. The monogenic signal is an isotropic extension of the analytic signal which preserves the core properties of the 1-D analytic signal that decomposes a signal into information about its structure (local phase) and energy (local amplitude). The analytic signal is generated using an isotropic vector valued odd filter known as the Riesz transform, which is a generalization of the Hilbert transform for higher dimensional signals. The spatial representations of these filters are as follow:

$$h_1(x, y) = \frac{-x}{2\pi(x^2 + y^2)^{3/2}} \quad (3)$$

$$h_2(x, y) = \frac{-y}{2\pi(x^2 + y^2)^{3/2}} \quad (4)$$

In practice, the image $I(x, y)$ is first convolved with an even isotropic band-pass filter $b(x, y)$ that produces the even component of the monogenic signal:

$$\text{even}(x, y) = I_b(x, y) = b(x, y) * I(x, y) \quad (5)$$

The bandpassed image $I_b(x, y)$ is then filtered with the Riesz filter to produce the odd components:

$$\text{odd}_1(x, y) = h_1(x, y) * I_b(x, y)$$

$$\text{odd}_2(x, y) = h_2(x, y) * I_b(x, y)$$

$$\text{odd}(x, y) = \sqrt{(\text{odd}_1(x, y))^2 + (\text{odd}_2(x, y))^2} \quad (6)$$

The local phase $\varphi(x, y)$ of the image $I(x, y)$ is produced through the following definition:

$$\varphi(x, y) = \arctan\left(\frac{\text{even}(x, y)}{\text{odd}(x, y)}\right) \quad (7)$$

The multi-scale local phase image $\varphi_{MS}(x, y)$ is formed by averaging the local phase computed at different scales:

$$\varphi_{MS}(x, y) = \frac{1}{N} \sum_s \varphi_s(x, y) \quad (8)$$

where s represents the scales and N is the total number of scales.

For the band-pass filter, we used a log-Gabor filter because it allows arbitrarily large bandwidth zero DC filters to be constructed. In 1-D, the log-Gabor filter has a transfer function of the form:

$$G(\omega) = \exp\left(-\frac{\log^2(\omega/k)}{2\log^2(\sigma_\omega)}\right)$$

where k is the centre frequency of the filter (which is inversely related to the scale of the filter), and $0 < \sigma_\omega < 1$ is related to the spread of the frequency spectrum in a logarithmic function. The following parameters produced the best empirical results: $\sigma_\omega = 0.50$ and 3 scales of filter wavelength: [250 150 50] pixels. Figure 1 shows the effect of using single-scale filter and multi-scale filter in producing the local phase (LP) image.




Figure 1: Example of local phase images produced by using different scales of filter

The stomach (*solid arrow*) and the umbilical vein (*dashed arrow*) are shown to have a better identification and separation from the surroundings in (e) the multi-scale local phase image as compared to (b-d) other single-scale local phase images.

3 Haar features

Features derived from Haar wavelets have been commonly used for representing the information in the image region. There are two motivations for using Haar features instead of pixel intensities directly. Firstly, features encode domain knowledge better than pixels for different types of visual patterns including objects in ultrasound images. The other reason is that a feature-based system can be much faster than a pixel-based system due to the use of the integral image. In their framework of object detection, Papageorgiou et al. [6] proposed the Haar-like features as a basis for image representation which was later

employed in a face detection technique [7]. The set of rectangle features provide a rich image representation which enables an effective learning process.




Figure 2: The unary, two-, three- and four- rectangles prototypes used in our implementation. The features are calculated by summing the value in black regions and subtracted with the totals in white regions (except for unary feature).

4 AdaBoost Learning Algorithm

For feature selections and classifier training, we used the AdaBoost [8] algorithm which is an established and proven effective method in object detection problems [7]. AdaBoost forms a strong hypothesis through linear combination of weak classifiers that are derived from the supplied pool of extracted features. The algorithm only requires the number of iterations (boosting) and automatically selects the most representative features during the training process.

The positive training samples for the classifier were cropped from the image regions that contain the anatomical object. Regions cropped from the background and the negative images (that do not contain the anatomical object) were used as negative training samples. We obtained 633 positive samples and 1032 negative samples for the stomach, and 448 positive samples and 965 negative samples for the umbilical vein training set.

Two different feature sets (*Intensity features* and *Local phase (LP) features*) were then extracted from the training dataset. Boosting process was performed separately on each feature set and also on the combination of both feature sets (*Intensity+LP*). The first five features chosen by the boosting algorithm are shown in Figure 3. Note that in both the stomach and the umbilical vein trained models, the local phase features were the first features selected by the learning algorithm and assigned with large classifier weights (α) which indicate high discriminating power of the features.




Figure 3: The first five features selected by AdaBoost for (a) the stomach and (b) the umbilical vein detection superimposed on the sample images. The greyscale and the coloured images indicate that the features selected are from the intensity feature set and local phase feature set, respectively.

For testing, exhaustive scan using the sliding window technique was performed (for all possible translations and a sparse set of scales) to find the anatomical object in the test image. The trained classifier was applied to all sub-windows within an image and takes the maximum of the classification score (α_T) as indication of the presence or absence of an object.

5 Data Acquisition

The fetal abdominal images used in this work were randomly selected from an ongoing clinical study database for fetal growth. The study involves normal singleton pregnancy from over 4000 healthy pregnant women screened at study entry to exclude risk factors for intrauterine growth restriction and overgrowth. Philips HD9 ultrasound machine with a 2-5MHz 2D probe are used by ultrasonographers trained to follow standardized procedures from study. The testing datasets consisting of 2384 images were labeled after consultation with trained sonographers.

6 Results

The performance of the detection methods trained with the three different feature sets were compared using ROC curves, as shown in Figure 4. Area under the curve (AUC) and balanced accuracy [$(\textit{specificity} + \textit{sensitivity})/2$] were summarized in Table 1.

Table 1: The performance of the detection methods using three different feature sets.

Feature sets	Area under the curve (AUC)		Balanced Accuracy (%)	
	Stomach	Umbilical Vein	Stomach	Umbilical Vein
Intensity	0.80	0.57	78.94	62.80
LP	0.66	0.48	67.40	56.18
Intensity + LP	0.83	0.63	80.14	65.16

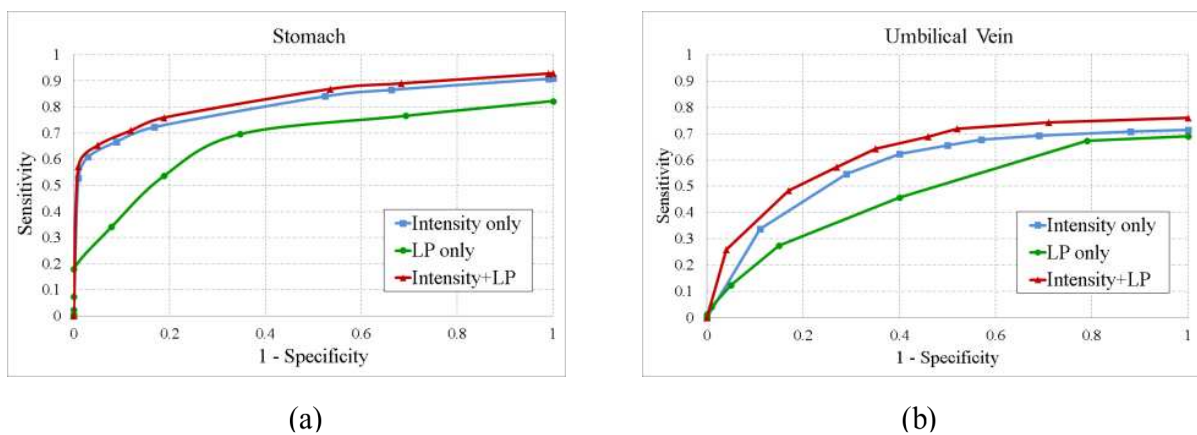


Figure 4: ROC plots for the detection of (a) the stomach and (b) the umbilical vein.

In comparison to the intensity-based features method, the proposed approach (*Intensity + LP* feature sets) achieved an increase of 1.20% and 2.36% in the accuracy of the stomach and the umbilical vein detection, respectively. These percentages translate to an increase of 100 true positive (TP) detections of stomach but with a decrease of 2 true

negative (TN) cases. For umbilical vein, it accounts for an increase of 62 TP and 2 TN detections.

7 Conclusions

In this work we have introduced a feature sets derived from multi-scale local phase images and intensity images into the machine learning framework for detecting anatomical landmarks (stomach and umbilical vein) in fetal abdominal ultrasound images. Quantitative results showed improved performance in comparison to the existing approach. This enhancement will in turn enable a more accurate quality assessment of fetal biometric ultrasound scans for improving the accuracy and the reproducibility of the fetal biometric measurements. This work is being extended to utilize other information available from local phase values (such as feature symmetry and asymmetry) and also the detection of objects in other fetal biometry scan areas.

References

- [1] L. S. Chitty, D. G. Altman, A. Henderson, and S. Campbell, "Charts of fetal size: 3. Abdominal measurements," *British Journal of Obstetrics and Gynaecology*, vol. 101, pp. 125-131, 1994.
- [2] B. Rahmatullah, I. Sarris, A. Papageorghiou, and J. A. Noble, "Quality control of fetal ultrasound images: Detection of abdomen anatomical landmarks using AdaBoost," in *2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, Chicago, IL, 2011, pp. 6-9.
- [3] V. Grau and J. A. Noble, "Adaptive multiscale ultrasound compounding using phase information," vol. 3749 LNCS, ed. Palm Springs, CA, 2005, pp. 589-596.
- [4] M. Mulet-Parada and J. A. Noble, "2D+T acoustic boundary detection in echocardiography," in *Medical Image Analysis*, 2000, pp. 21-30.
- [5] M. Felsberg and G. Sommer, "The monogenic signal," *IEEE Transactions on Signal Processing*, vol. 49, pp. 3136-3144, 2001.
- [6] M. Oren, C. Papageorgiou, P. Sinha, E. Osuna, and T. Poggio, "Pedestrian detection using wavelet templates," in *Computer Vision and Pattern Recognition (CVPR)*, 1997, pp. 193-199.
- [7] P. Viola and M. J. Jones, "Robust Real-Time Face Detection," *International Journal of Computer Vision*, vol. 57, pp. 137-154, 2004.
- [8] Y. Freund and R. E. Schapire, "A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting," *Journal of Computer and System Sciences*, vol. 55, pp. 119-139, 1997.