



Standard Plane Detection in 3D Fetal Ultrasound Using an Iterative Transformation Network

Yuanwei Li¹✉, Bishesh Khanal², Benjamin Hou¹, Amir Alansary¹, Juan J. Cerrolaza¹, Matthew Sinclair¹, Jacqueline Matthew², Chandni Gupta², Caroline Knight², Bernhard Kainz¹, and Daniel Rueckert¹

¹ Biomedical Image Analysis Group, Imperial College London, London, UK
yuanwei.li09@imperial.ac.uk

² School of Biomedical Engineering and Imaging Sciences,
King's College London, London, UK

Abstract. Standard scan plane detection in fetal brain ultrasound (US) forms a crucial step in the assessment of fetal development. In clinical settings, this is done by manually manoeuvring a 2D probe to the desired scan plane. With the advent of 3D US, the entire fetal brain volume containing these standard planes can be easily acquired. However, manual standard plane identification in 3D volume is labour-intensive and requires expert knowledge of fetal anatomy. We propose a new Iterative Transformation Network (ITN) for the automatic detection of standard planes in 3D volumes. ITN uses a convolutional neural network to learn the relationship between a 2D plane image and the transformation parameters required to move that plane towards the location/orientation of the standard plane in the 3D volume. During inference, the current plane image is passed iteratively to the network until it converges to the standard plane location. We explore the effect of using different transformation representations as regression outputs of ITN. Under a multi-task learning framework, we introduce additional classification probability outputs to the network to act as confidence measures for the regressed transformation parameters in order to further improve the localisation accuracy. When evaluated on 72 US volumes of fetal brain, our method achieves an error of 3.83 mm/12.7° and 3.80 mm/12.6° for the transventricular and transcerebellar planes respectively and takes 0.46 s per plane.

1 Introduction

Obstetric ultrasound (US) is conducted as a routine screening examination between 18–24 weeks of gestation. US imaging of the fetal head enables clinicians to assess fetal brain development and detect growth abnormalities. This

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-00928-1_45) contains supplementary material, which is available to authorized users.

requires the careful selection of standard scan planes such as the transventricular (TV) and transcerebellar (TC) plane that contain key anatomical structures [6]. However, it is challenging and time-consuming even for experienced sonographers to manually navigate a 2D US probe to find the correct standard plane. The task is highly operator-dependent and requires a great amount of expertise. With the advent of 3D fetal US, a volume of the entire fetal brain can be acquired quickly with little training. But the problem of locating diagnostically required standard planes for biometric measurements remains. There is a strong need to develop automatic methods for 2D standard plane extraction from 3D volumes to improve clinical workflow efficiency.

Related work: Recently, deep learning approaches have shown successes in many medical image analysis applications. Several works have applied deep learning techniques to standard plane detection in fetal US [1–3, 7]. Baumgartner *et al.* [1] use a convolutional neural network (CNN) for categorisation of 13 fetal standard views. Chen *et al.* [3] adopt a CNN-based image classification approach for detecting fetal abdominal standard planes, which they later combined with a recurrent neural network (RNN) that takes into account temporal information [2]. However, these methods identify standard planes from 2D US videos and not 3D volumes. Ryou *et al.* [7] attempt to detect fetal head and abdominal planes from 3D fetal US by breaking down the 3D volume into a stack of 2D slices which are then classified as head or abdomen using a CNN.

Plane detection is considered an image classification problem in the above works. In contrast, we approach the plane detection problem by regressing rigid transformation parameters that define the plane position and orientation. There are several works on using CNN to predict transformations. Kendall *et al.* [5] introduce PoseNet for regressing 6-DoF camera pose from RGB image with a loss function that uses quaternions to represent rotation. Hou *et al.* [4] propose SVRNet for predicting transformation from 2D image to 3D space and use anchor points as a new representation for rigid transformations. These works predict absolute transformation with respect to a known reference coordinate system with one pass of CNN. Our work is different as we use an iterative approach with multiple passes of CNN to predict relative transformation with respect to current plane coordinates, which change at each iteration. Relative transformation is used as our 3D volumes are not aligned to a reference coordinate system.

Contributions: In this paper, we propose the Iterative Transformation Network (ITN) that uses a CNN to detect standard planes in 3D fetal US. The network learns a mapping between a 2D plane and the transformation required to move that plane towards the standard plane within a 3D volume. Our contributions are threefold: **(1)** ITN is a general deep learning framework built for 2D plane detection in 3D volumes. The iterative approach regresses transformations that bring the plane closer to the standard plane. This reduces computation cost as ITN selectively samples only a few planes in the 3D volume unlike classification-based methods that require dense sampling [1–3, 7]. **(2)** We study the effect on plane detection accuracy using different transformation representations (quaternions, Euler angles, rotation matrix, anchor points) as CNN regression outputs. **(3)** We




Fig. 1. (a) Overall plane detection framework using ITN. (b) Composition of transformations. Red: GT plane. Blue: Arbitrary plane. Black: Identity plane.

improve ITN performance by incorporating additional classification probability outputs as confidence measures of the regressed transformation parameters. At inference, the classification probabilities are used as confidence scores to yield more accurate localisation. During training, regression and classification outputs are learned in a multi-task learning framework, which improves the generalisation ability of the model and prevents overfitting.

2 Method

Overall Framework: Fig. 1a presents the overall ITN framework for plane detection. Given a 3D volume V , the goal is to find the ground truth (GT) standard plane (red). Starting with a random plane initialisation (blue), the 2D image of the plane is extracted and input to a CNN which then predicts a 3D transformation ΔT that will move the plane to a new position closer to the GT plane. The image extracted at the new plane location is then passed to the CNN and the process is repeated until the plane reaches the GT plane.

Composition of Transformations: Transformation is defined with respect to a reference coordinate system. In Fig. 1b, we define an identity plane (black) with origin at the volume centre. T and T^{GT} are defined in the coordinate system of the identity plane and they move the identity plane to the arbitrary plane (blue) and GT plane (red) respectively. ΔT^{GT} is defined in the coordinate system of the arbitrary plane and ΔT^{GT} moves the arbitrary plane to the GT plane. Note that our ITN predicts ΔT^{GT} which is a relative transformation from the point of view of the current plane, and *not* from the identity plane. We compute these transformations from each other using $T^{GT} = T \oplus \Delta T^{GT}$ and $\Delta T^{GT} = T^{GT} \ominus T$ where \oplus and \ominus are the composition and inverse composition operators respectively. The computations defined by the operators are dependent on the choice of the transformation representation.

Network Training: During training, an arbitrary plane is randomly sampled from a volume V by applying a random transformation T to the identity plane. The corresponding 2D plane image X is then extracted. We define $X = I(V, T, s)$ where $I(\cdot)$ is the plane extraction function and s is the length of the square plane. We sample T such that the plane centre falls in the middle 60% of V and the rotation of the plane is within an angle of $\pm 45^\circ$ about each coordinate axis.

Algorithm 1. Iterative Inference of Transformation

```

1: procedure PLANE TRANSFORMATION( $V, s, N$ )
2:   Initialise random plane with  $T_1$ 
3:   for  $i = 1$  to  $N$  do
4:      $X_i \leftarrow I(V, T_i, s)$  ▷ Sample plane image
5:      $\Delta T \leftarrow \text{CNN}(X_i)$  ▷ CNN predicts relative transformation
6:      $T_{i+1} \leftarrow T_i \oplus \Delta T$  ▷ Update plane position
7:   return  $T_N$ 

```

Table 1. Representations of rigid transformations and their loss functions.

Representation (Parameter count)	Loss function
Translation \mathbf{t} (3) + Quaternion \mathbf{q} (4)	$L = \alpha \ \mathbf{t}^{GT} - \mathbf{t}\ _2^2 + \beta \left\ \mathbf{q}^{GT} - \frac{\mathbf{q}}{\ \mathbf{q}\ } \right\ _2^2$
Translation \mathbf{t} (3) + Euler angles $\boldsymbol{\theta}$ (3)	$L = \alpha \ \mathbf{t}^{GT} - \mathbf{t}\ _2^2 + \beta \ \boldsymbol{\theta}^{GT} - \boldsymbol{\theta}\ _2^2$
Translation \mathbf{t} (3) + Rotation matrix \mathbf{R} (9)	$L = \alpha \ \mathbf{t}^{GT} - \mathbf{t}\ _2^2 + \beta \ \mathbf{R}^{GT} - \mathbf{R}\ _2^2$
Anchor points $(\mathbf{A}_1, \mathbf{A}_2, \mathbf{A}_3)$ (9)	$L = \sum_{i=1}^3 \ \mathbf{A}_i^{GT} - \mathbf{A}_i\ _2^2$

This avoids sampling of planes at the edges of the volume where there is no informative image data due to regions falling outside of the US imaging cone.

A training sample is represented by $(X, \Delta T^{GT})$ and the training loss function can be formulated as the $L2$ norm of the error between the GT and predicted transformation parameters: $L = \|\Delta T^{GT} - \Delta T\|_2^2$

Network Inference: Algorithm 1 summarises the steps during network inference to detect a plane. The iterative approach gives rough estimates of the plane in the first few iterations and subsequently makes smaller and more accurate refinements. This coarse-to-fine adjustment improves accuracy and is less susceptible to different initialisations. To improve accuracy and convergence, we repeat Algorithm 1 with 5 random plane initialisations per volume and average their final transformations T_N after N iterations.

Transformation Representations: In ITN, plane transformation ΔT is rigid, comprising only translation and rotation. We explore the effect of using different transformation representations as the CNN regression outputs (Table 1) since there are few comparative studies that investigate this on deep networks. The first three representations explicitly separate translation and rotation in which rotation is represented by quaternions, Euler angles and rotation matrix respectively. α and β are weightings given to the translation and rotation losses. Specifically, anchor points [4] are defined as the coordinates of three fixed points on the plane (we use: centre, bottom-left and bottom-right corner). The points uniquely and jointly represent any translation and rotation in 3D space. During inference, the predicted values of certain representations need to be constrained to give valid rotation. For instance, quaternions need to be normalised to unit quaternions and rotation matrices need to be orthogonalised. Anchor points need to be converted to valid rotation matrices as described in [4].

Algorithm 2. Compute relative transformation ΔT

```

1: procedure COMPUTE_TRANSFORM( $\mathbf{t}, \mathbf{q}, \mathbf{P}, \mathbf{Q}$ )
2:    $\mathbf{t}_{new} = \begin{pmatrix} \max(P_{c_1^+}, P_{c_1^-})t_1 \\ \max(P_{c_2^+}, P_{c_2^-})t_2 \\ \max(P_{c_3^+}, P_{c_3^-})t_3 \end{pmatrix}$  ▷ Compute weighted translation
3:    $Q_{max} = \max(\mathbf{Q})$  ▷ Compute weighted rotation
4:   if  $Q_{max} = Q_{k_x^+}$  OR  $Q_{k_x^-}$  then
5:     Convert  $\mathbf{q}$  to Euler angles  $(\theta_x, \theta_y, \theta_z)$  using convention ‘xyz’
6:      $\mathbf{r}_{new} \leftarrow$  Rotation about x-axis with magnitude  $Q_{max}\theta_x$ 
7:   else if  $Q_{max} = Q_{k_y^+}$  OR  $Q_{k_y^-}$  then
8:     Convert  $\mathbf{q}$  to Euler angles  $(\theta_x, \theta_y, \theta_z)$  using convention ‘yxz’
9:      $\mathbf{r}_{new} \leftarrow$  Rotation about y-axis with magnitude  $Q_{max}\theta_y$ 
10:  else if  $Q_{max} = Q_{k_z^+}$  OR  $Q_{k_z^-}$  then
11:    Convert  $\mathbf{q}$  to Euler angles  $(\theta_x, \theta_y, \theta_z)$  using convention ‘zxy’
12:     $\mathbf{r}_{new} \leftarrow$  Rotation about z-axis with magnitude  $Q_{max}\theta_z$ 
13:   $\Delta T \leftarrow (\mathbf{t}_{new}, \mathbf{r}_{new})$ 
14:  return  $\Delta T$ 

```

Classification Probability as Confidence Measure: We further extend our ITN by incorporating classification probability as a confidence measure for the regressed values of translation and rotation. The method can be applied to any transformation representation but we use quaternions since it yields the best results. In addition to the regression outputs \mathbf{t} and \mathbf{q} , the CNN also predicts two classification probability outputs \mathbf{P} and \mathbf{Q} for translation and rotation respectively. We divide translation into 6 discrete classification categories: positive and negative translation along each coordinate axis. Denoting c as the translation classification label, we have $c \in \{c_1^+, c_1^-, c_2^+, c_2^-, c_3^+, c_3^-\}$ where c_1^+ is the category representing translation along the positive x-axis. \mathbf{P} is then a 6-element vector giving the probability of translation along each axis direction. Similarly, we divide rotation into 6 categories: clockwise and counter-clockwise rotation about each coordinate axis. Denoting k as the rotation classification label, we have $k \in \{k_1^+, k_1^-, k_2^+, k_2^-, k_3^+, k_3^-\}$ where k_1^+ is the category representing clockwise rotation about the x-axis. \mathbf{Q} is then a 6-element vector giving the probability of rotation about each axis.

A training sample is represented by $(X, \mathbf{t}^{GT}, \mathbf{q}^{GT}, c^{GT}, k^{GT})$. c^{GT} gives the coordinate axis along which the current plane centre has the furthest absolute distance from the GT plane centre. Similarly, k^{GT} gives the coordinate axis about which the current plane will rotate the most to reach the GT plane. Appendix A derives the computations of c^{GT} and k^{GT} during training. The overall training loss function can then be written as:

$$L = \alpha \|\mathbf{t}^{GT} - \mathbf{t}\|_2^2 + \beta \left\| \mathbf{q}^{GT} - \frac{\mathbf{q}}{\|\mathbf{q}\|} \right\|_2^2 - \gamma \log P_{c^{GT}} - \delta \log Q_{k^{GT}} \quad (1)$$

The first and second terms are the $L2$ losses for translation and rotation regression while the third and fourth terms are the cross-entropy losses for translation and rotation classification. α , β , γ and δ are weights given to the losses.

During inference, the CNN outputs \mathbf{t} , \mathbf{q} , \mathbf{P} and \mathbf{Q} are combined to compute the relative transformation ΔT (Algorithm 2). For translation, each component of the regressed translation \mathbf{t} is weighted by the corresponding probabilities in the vector \mathbf{P} . For rotation, we only rotate the plane about the most confident rotation axis as predicted by \mathbf{Q} . In order to determine the magnitude of that rotation, the regressed quaternion \mathbf{q} needs to be broken down into Euler angles using the appropriate convention in order to determine the rotation angle about that most confident rotation axis. An Euler angle representation using convention ‘xyz’ means a rotation about x-axis first followed by y-axis and finally z-axis. Hence, \mathbf{P} and \mathbf{Q} are used as confidence weighting for \mathbf{t} and \mathbf{q} , allowing the plane to translate and rotate to a greater extent along the more confident axis.

Network Architecture: ITN utilises a multi-task learning framework for predictions of multiple outputs. The architecture differs according to the number of outputs that the CNN predicts. All our networks comprise 5 convolution layers, each followed by a max-pooling layer. These layers contain shared features for all outputs. After the 5th pooling layer, the network branches into fully-connected layers to learn the specific features for each output. Details of all network architectures are described in Appendix B.

3 Experiments and Results

Data and Experiments: ITN is evaluated on 3D US volumes of fetal brain from 72 subjects. For each volume, TV and TC standard planes are manually selected by a clinical expert. 70% of the dataset is randomly selected for training and the rest 30% used for testing. All volumes are processed to be isotropic with mean dimensions of $324 \times 207 \times 279$ voxels. ITN is implemented using Tensorflow running on a machine with Intel Xeon CPU E5-1630 at 3.70 GHz and one NVIDIA Titan Xp 12GB GPU. We set plane size $s=225$, $N=10$ and $\alpha=\beta=\gamma=\delta=1$. During training, we use a batch size of 64. Weights are initialised randomly from a distribution with zero mean and 0.1 standard deviation. Optimisation is carried out for 100,000 iterations using the Adam algorithm with learning rate=0.001, $\beta_1=0.9$ and $\beta_2=0.999$. The predicted plane is evaluated against the GT using distance between the plane centres (δx) and rotation angle between the planes ($\delta \theta$). Image similarity of the planes is also measured using peak signal-to-noise ratio (PSNR) and structural similarity (SSIM).

Results: Table 2 compares the plane detection results when different transformation representations are used by ITN. In general, there is little difference in the translation error. This is because all translation representations are the same, which use the three Cartesian axes except for anchor points which have slightly greater translation error. The rotation errors on TC plane suggest that quaternions are a good representation. Rotation matrices and anchor points

Table 2. Evaluation of ITN with different transformation representations for standard plane detection. Results presented as (Mean \pm Standard Deviation).

CNN outputs	TV plane				TC plane			
	δx (mm)	$\delta\theta$ ($^\circ$)	PSNR	SSIM	δx (mm)	$\delta\theta$ ($^\circ$)	PSNR	SSIM
t, q	6.29 \pm 5.33	17.0 \pm 12.0	15.3 \pm 2.0	0.375 \pm 0.081	6.23 \pm 6.99	14.9 \pm 7.5	15.5 \pm 2.1	0.383 \pm 0.100
t, θ	5.69 \pm 5.85	17.0 \pm 8.5	15.2 \pm 1.7	0.372 \pm 0.084	7.13 \pm 9.00	16.0 \pm 5.9	14.6 \pm 2.4	0.357 \pm 0.119
t, R	5.79 \pm 6.10	17.7 \pm 11.6	15.8 \pm 1.9	0.389 \pm 0.091	6.39 \pm 7.39	17.3 \pm 15.4	15.5 \pm 2.4	0.385 \pm 0.118
A_1, A_2, A_3	6.64 \pm 8.66	17.0 \pm 10.4	15.9 \pm 2.4	0.399 \pm 0.099	7.88 \pm 10.0	16.3 \pm 12.6	15.0 \pm 2.7	0.351 \pm 0.124

over-parameterise rotation and can make network learning more difficult with greater degree of freedom. Since these parameters are not constrained, it is also harder to convert them back into valid rotations during inference. Quaternions have fewer parameters and slightly-off quaternion can still be easily normalised to give valid rotation. Compared to Euler angles, quaternions avoid the problem of gimbal lock. For TV plane, there is little difference in rotation error. This is because sonographers use the TV plane as a visual reference when acquiring 3D volumes. This causes the TV plane to lie roughly in the central plane of the volume with lower rotation variances, thus making the choice of rotation representation less important. Table 3 compares the performance of ITN with/without classification probability outputs. Given a baseline model (M1) that only has regression outputs t, q , the addition of classification probabilities P, Q improves the translation and rotation accuracy respectively (M2-M4). The classification probabilities act as confidence weights for the regression outputs to improve plane detection accuracy. Furthermore, the classification and regression outputs are trained in a multi-task learning fashion, which allows feature sharing and enables more generic features to be learned, thus preventing model overfitting. M1-M4 use one plane image as CNN input. We further improve our results by using three orthogonal plane images instead as this provides more information about the 3D volume (M4+). M4 and M4+ take 0.46 s and 1.35 s to predict one plane per volume. The supplementary material provides videos showing the update of a randomly initialised plane and its extracted image through 10 inference iterations.

Figure 2 shows a visual comparison between the GT planes and the planes predicted by M4. To evaluate the clinical relevance of the predicted planes, a

Table 3. Evaluation of ITN with/without confidence probability for standard plane detection. Results presented as (Mean \pm Standard Deviation).

CNN outputs	TV plane				TC plane			
	δx (mm)	$\delta\theta$ ($^\circ$)	PSNR	SSIM	δx (mm)	$\delta\theta$ ($^\circ$)	PSNR	SSIM
M1: t, q	6.29 \pm 5.33	17.0 \pm 12.0	15.3 \pm 2.0	0.375 \pm 0.081	6.23 \pm 6.99	14.9 \pm 7.5	15.5 \pm 2.1	0.383 \pm 0.100
M2: t, q, P	5.14 \pm 5.37	16.8 \pm 9.9	16.0 \pm 2.1	0.408 \pm 0.092	5.12 \pm 5.50	13.9 \pm 7.1	15.8 \pm 2.2	0.393 \pm 0.115
M3: t, q, Q	6.07 \pm 6.32	14.0 \pm 8.1	15.7 \pm 2.5	0.399 \pm 0.108	7.66 \pm 7.14	12.7 \pm 6.0	15.5 \pm 3.0	0.386 \pm 0.123
M4: t, q, P, Q	3.83 \pm 2.10	12.7 \pm 7.7	16.4 \pm 1.9	0.419\pm0.092	3.80 \pm 1.85	12.6 \pm 6.1	16.5 \pm 2.1	0.407 \pm 0.110
M4+: t, q, P, Q	3.49\pm1.81	10.7\pm5.7	16.6\pm1.8	0.413 \pm 0.082	3.39\pm2.13	11.4\pm6.3	16.8\pm2.1	0.437\pm0.110

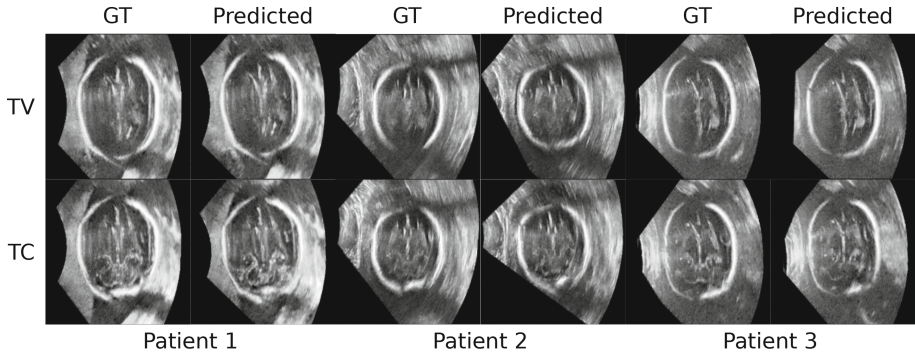


Fig. 2. Visualisation of GT planes and planes predicted by M4.

clinical expert manually measures the head circumference (HC) on both the predicted and GT planes and computes the standard deviation of the measurement error to be 1.05 mm (TV) and 1.25 mm (TC). This is similar to the intraobserver variability of 2.65 mm reported for HC measurements on TC plane [8]. Thus, accurate biometrics can be extracted from our predicted planes.

4 Conclusion

We presented ITN, a new approach for standard plane detection in 3D fetal US by using a CNN to regress rigid transformation iteratively. We compare the use of different transformation representations and show quaternions to be a good representation for iterative pose estimation. Additional classification probabilities are learned via multi-task learning which act as confidence weights for the regressed transformation parameters to improve plane detection accuracy. As future work, we are evaluating ITN on other plane detection tasks (*eg.* view plane selection in cardiac MRI). It is also worthwhile to explore new transformation representations and extend ITN to simultaneous detection of multiple planes.

Acknowledgments. Supported by the Wellcome Trust IEH Award [102431]. The authors thank Nvidia Corporation for the donation of a Titan Xp GPU.

References

1. Baumgartner, C.F., et al.: Sononet: real-time detection and localisation of fetal standard scan planes in freehand ultrasound. *IEEE TMI* **36**(11), 2204–2215 (2017)
2. Chen, H., et al.: Automatic fetal ultrasound standard plane detection using knowledge transferred recurrent neural networks. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) *MICCAI 2015*. LNCS, vol. 9349, pp. 507–514. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24553-9_62

3. Chen, H., Ni, D., Qin, J., Li, S., Yang, X., Wang, T., Heng, P.A.: Standard plane localization in fetal ultrasound via domain transferred deep neural networks. *IEEE J. Biomed. Health Inf.* **19**(5), 1627–1636 (2015)
4. Hou, B., et al.: Predicting slice-to-volume transformation in presence of arbitrary subject motion. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) *MICCAI 2017*. LNCS, vol. 10434, pp. 296–304. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66185-8_34
5. Kendall, A., Grimes, M., Cipolla, R.: PoseNet: a convolutional network for real-time 6-dof camera relocalization. In: *ICCV 2015*, pp. 2938–2946. IEEE (2015)
6. NHS: Fetal anomaly screening programme: programme handbook June 2015. Public Health England (2015)
7. Ryou, H., Yaqub, M., Cavallaro, A., Roseman, F., Papageorghiou, A., Noble, J.A.: Automated 3D ultrasound biometry planes extraction for first trimester fetal assessment. In: Wang, L., Adeli, E., Wang, Q., Shi, Y., Suk, H.-I. (eds.) *MLMI 2016*. LNCS, vol. 10019, pp. 196–204. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-47157-0_24
8. Sarris, I., et al.: Intra- and interobserver variability in fetal ultrasound measurements. *Ultrasound Obstet. Gynecol.* **39**(3), 266–273 (2012)