Toward point-of-care ultrasound estimation of fetal gestational age from the trans-cerebellar diameter using CNN-based ultrasound image analysis

Mohammad A. Maraci
Mohammad Yaqub
Rachel Craik
Sridevi Beriwal
Alice Self
Peter von Dadelszen
Aris Papageorghiou
J. Alison Noble

Toward point-of-care ultrasound estimation of fetal gestational age from the trans-cerebellar diameter using CNN-based ultrasound image analysis

Mohammad A. Maraci,^a Mohammad Yaqub,^a Rachel Craik,^{b,c} Sridevi Beriwal,^b Alice Self,^b Peter von Dadelszen,^c Aris Papageorgiou,^b and J. Alison Noble^{a,*}

^aUniversity of Oxford, Institute of Biomedical Engineering, Department of Engineering Science, Oxford, United Kingdom

^bUniversity of Oxford, Nuffield Department of Women's and Reproductive Health, Oxford, United Kingdom

^cKing's College London, Department of Women and Children's Health, London, United Kingdom

Abstract. Obstetric ultrasound is a fundamental ingredient of modern prenatal care with many applications including accurate dating of a pregnancy, identifying pregnancy-related complications, and diagnosis of fetal abnormalities. However, despite its many benefits, two factors currently prevent wide-scale uptake of this technology for point-of-care clinical decision-making in low- and middle-income country (LMIC) settings. First, there is a steep learning curve for scan proficiency, and second, there has been a lack of easy-to-use, affordable, and portable ultrasound devices. We introduce a framework toward addressing these barriers, enabled by recent advances in machine learning applied to medical imaging. The framework is designed to be realizable as a point-of-care ultrasound (POCUS) solution with an affordable wireless ultrasound probe, a smartphone or tablet, and automated machine-learning-based image processing. Specifically, we propose a machine-learning-based algorithm pipeline designed to automatically estimate the gestational age of a fetus from a short fetal ultrasound scan. We present proof-of-concept evaluation of accuracy of the key image analysis algorithms for automatic head transcerebellar plane detection, automatic transcerebellar diameter measurement, and estimation of gestational age on conventional ultrasound data simulating the POCUS task and discuss next steps toward translation via a first application on clinical ultrasound video from a low-cost ultrasound probe. © 2020 Society of Photo-Optical Instrumentation Engineers (SPIE) [DOI: 10.1117/1.JMI.7.1.014501]

Keywords: point-of-care ultrasound; prenatal health; machine learning; global health; gestational age.

Paper 19026RR received Jan. 28, 2019; accepted for publication Dec. 5, 2019; published online Jan. 13, 2020.

1 Introduction

According to the World Health Organization, 99% of life-ending maternal complications in pregnancy occur in low- and middle-income countries (LMICs).¹ Further, optimizing maternal health in pregnancy is designated to be one of the global challenges in 21st century medicine as diagnosing pregnancy-related conditions can be complex and is compounded by infrastructural constraints imposed by inadequate accessibility, insufficient expertise, and requirement of capital equipment. Ultrasound is the recommended imaging modality for monitoring fetal and maternal health during pregnancy, as it is highly portable and economic, facilitates real-time data acquisition and visualization, and also benefits from an acceptable safety profile for patients and sonographers alike.² Recently, the use of medical ultrasound in limited-resource settings has increased,^{3–6} in part attributable to the increased portability and affordability of this technology.

Demand for point-of-care ultrasound (POCUS) solutions in first response scenarios and remote locations, the emergency department and at the patient's bedside, have driven recent advances in ultrasound technology portability and affordability. Easily transportable and portable handheld ultrasound systems can turn either a consumer-grade smartphone or tablet into a portable ultrasound system, thus providing an affordable POCUS solution. Empowering nonexperts or minimally trained

healthcare professionals to use such systems is not straightforward as the quality of ultrasound image acquisition is depended on user skill. Therefore, an important emerging research area is automated image analysis to support nonexpert users in performing ultrasound-based diagnostic decisions. This paper contributes to this literature.

Specifically, in this paper, we present an image analysis framework suitable for processing ultrasound video in real-time for automatic detection and measurement of the transcerebellar diameter (TCD) in the fetal brain, which enables the estimation of fetal gestational age (GA). Fetal GA is a key pregnancy biomarker that underlies much of clinical decision-making and is often very inaccurate in LMICs.

1.1 Gestational Age Estimation

Accurate estimation of GA based on sonographic fetal parameters forms an integral part of modern prenatal care. In addition to defining the expected delivery date, it is also essential for screening of anatomic and growth abnormalities and disturbances, and it allows appropriate scheduling of antenatal care. At a population level, reliable estimates of GA are important to calculate rates of preterm birth and small-for-gestational age (SGA) at birth. Currently, estimates of preterm birth and SGA rates are mere approximations of true rates, particularly in geographical regions at greatest risk of these conditions.^{7,8}

*Address all Correspondence to J. Alison Noble, E-mail: alison.noble@eng.ox.ac.uk

Traditionally, GA is estimated using the first day of the last menstrual period, which assumes that ovulation occurs on day 14 day of a regular 28-day menstrual cycle. Irregular menses, unknown or uncertain dates, oral contraceptive use, or recent pregnancy or breastfeeding may all influence the accuracy of this method.^{9,10} In high-income countries (HICs), in such a case, ultrasound measurement of the fetal crown rump length in the first trimester is recommended,¹¹ as this is said to be more accurate than dating in late pregnancy, where growth disturbances result in underestimation of GA for an abnormally small fetus and an overestimation of GA for a macrosomic fetus. In the second and third trimesters, GA is typically estimated from fetal biparietal diameter, head circumference (HC), abdominal circumference (AC), or femur length (FL). As the TCD is least likely to be affected by fetal growth disturbances, it has been shown to be a single fetal biometric parameter that allows accurate prediction of GA in both singleton and twin gestations.¹² Therefore, the TCD was the parameter chosen for automated GA estimation in our work. To our knowledge, this is the first paper to present an automated approach to TCD-based GA estimation from ultrasound images and video.

1.2 Ultrasound in LMICs

A considerable challenge in LMICs is that women frequently present to a healthcare professional late in pregnancy (later than 20 weeks) and often do not, with confidence, know how many weeks pregnant they are. Therefore, accurate ultrasound-based estimation of GA is essential for implementation of evidence-based interventions, such as administration of corticosteroids for fetal lung maturation or induction of labor due to postmaturity.

Obstetrical ultrasound is acceptable to pregnant women and health care workers in many LMIC settings, although, in some settings, it is acknowledged that there is resistance to its use as ultrasound imaging can be used for fetal sex determination.¹³ Obstetrical ultrasound training programs reduce measurement variability¹⁴ and improve image quality.¹⁵ Observational studies have reported that measurement reproducibility of trained healthcare workers in limited-resource settings is comparable with that of qualified sonographers from HICs for certain basic diagnostic tasks.¹⁶ However, other work reports that reliable ultrasound-based measurement requires significant pre- and in-service training and is difficult to scale up.¹³ In response to this, DeStigter et al.⁶ proposed a care model by which minimally trained primary care workers followed a simple scanning protocol for acquiring freehand ultrasound sweeps (cine loops) over the abdomen [called an obstetric sweep protocol (OSP)¹⁷], and data were transferred to a center where experts reviewed the scans. The OSP essentially allows a healthcare professional to implement the ISUOG basic obstetric ultrasound checklist (Table 1 in Ref. 18). A recent comparison of three low-cost ultrasound devices for manual estimation of GA using both a standard plane and the OSP concluded that manual analysis using low-cost devices and the OSP gave good GA estimates from measurements of HC and AC but not FL;¹⁷ that paper was the first to report on GA estimation accuracy from ultrasound video sweeps that might be acquired by a minimally trained healthcare professional rather than assuming that the sonographer has the skill to find the correct biometry plane. However, the selection of the OSP-derived biometry plane and manual GA estimation were performed by an experienced sonographer.

The high prevalence of growth restriction in LMICs makes established methods of late GA assessment, such as measurement

of fetal HC, inaccurate. A proposed alternative has been measurement of the fetal TCD. The fetal TCD is a good measure for GA assessment over a wider range of GA and is relatively protected from the influence of fetal growth restriction in the absence of associated anomalies.¹² Thus, a TCD-based GA estimation device could be very useful in LMICs.

1.3 Automated Fetal Ultrasound Image Analysis

We restrict the discussion to relevant literature in two-dimensional (2-D) ultrasound image analysis as three-dimensional (3-D) ultrasound is currently too expensive to be considered for GA estimation in LMICs. For a discussion of the latter and a method for automatic 3-D ultrasound-based GA estimation, we refer the reader to Ref. 19.

There is growing academic literature on automated fetal ultrasound image analysis from 2-D ultrasound and video, for instance Refs 20–31. Much of this recent literature is focused on automated estimation of biometry²⁰ and automatic recognition of standard imaging planes for the second-trimester anomaly scan from ultrasound video frames.^{21–26} Other research has considered automatic ultrasound video partitioning and automatic video frame description to label key anatomical structures in the fetal heart.²⁷ Early methods were based on frame-by-frame analysis but recently a number of papers have demonstrated that superior plane recognition accuracy can be achieved with spatiotemporal models built with recurrent neural network and long short-term memory-based architectures.^{25,26} These make use of frame-to-frame coherence of local structures within the models to overcome ambiguity that can occur in single-frame analysis. In all of the previously cited papers, methods were motivated by the clinical need to support second-trimester obstetrics screening and fetal growth monitoring in HICs.

There is limited literature on automated fetal ultrasound image analysis aimed at clinical needs in LMIC settings.^{28–31} Reference 28 was the first paper to consider automation of aspects of the ISUOG basic obstetrical ultrasound scan.¹⁸ Specifically, Ref. 28 presented an approach to detect fetal presentation, that the heart is beating (fetal viability), and to estimate biometry from a single linear ultrasound sweep over the maternal abdomen. Gao et al. further developed that work to propose spatiotemporal models for linear sweep partitioning, where the most notable improvement was for fetal heart detection.^{29,30} More recently, van der Heuvel et al.³¹ proposed deep learning-based methods for automated estimation of fetal presentation, twin detection, and HC-based GA estimation from the OSP. While this is the most closely related work to our own, there are key differences: van der Heuvel et al. and others^{28–31} used machine-learning-based ultrasound video partitioning for object detection. As the OSP is a whole abdomen scan, the fetal head position with respect to the sweep is unknown. Hence, any given image frame is not guaranteed to be a standardized transventricular (TV) plane to provide the correct biometry plane for HC-based GA estimation. Therefore, in Ref. 31, HC-based GA estimation involved empirical combination of HC estimates from a number of frames. It is difficult to judge from that paper how to generalize that strategy to other measurements.

To our knowledge, there is no previous publication on automated TCD-based GA estimation implemented within an automated machine-learning-based framework. In the proposed approach, we assume (1) a different ultrasound sweep protocol than Ref. 31 and (2) that a minimally trained healthcare professional can find the fetal head and acquire a short video sweep of

the head by moving the ultrasound probe from a starting position of approximately (but not exactly) at the TV position. Then, the automated image analysis problem is to find the standardized TC plane and then to estimate the TCD from this cine loop.

2 Methods

2.1 Experimental Data

For the proof-of-principle purposes of this paper, we used two exemplar datasets.

Dataset A: Our underlying assumption is that the healthcare professional can find the head and approximately locate the TV plane. Then, they move the probe around slowly (small rotations and translations) while acquiring a short cine loop that includes the TC plane. The image analysis challenge is to find the TC plane when many candidates are actually TV planes. The difference between the two imaging planes is illustrated schematically in Fig. 1.

At the time of this research, we did not have access to a large real-world video dataset so we simulated this case by using a dataset of 5000 TC and TV images. This dataset was divided into a 3000 image training dataset consisting of 1500 images of the TC plane as positive examples and 1500 images of the TV plane as negative examples. Subsequently, the validation set consisted of 1000 unseen data samples with 500 TC (positive) and 500 TV (negative) images. This dataset was used to evaluate the plane finder.

Dataset B: Image dataset B was used to develop and evaluate the automatic TCD measurement algorithm. 3736 retrospective TC images were obtained from a second-trimester fetal ultrasound screening database. 500 images were randomly selected for the training set and the TC structure manually annotated, and the remaining 3236 images were used in testing. The TCD had been manually measured by a sonographer for each image during the scanning session. The manually measured TCD value was written on the image. Optical character recognition was used to read the manually measured TCD (mm) from the image. This provided the TCD measurement ground truth. According to

the manual TCD measurements, the GA range for these images covered 16- to 26-weeks GA.

These images were in DICOM format and hence pixel size in millimeters was available. This meant that we could directly convert the automatic TCD from pixels to millimeters.

2.2 Acquisition and Automated Analysis Pipeline

A schematic summarizing the acquisition and automated analysis pipeline principle is shown in Fig. 2. We briefly explain the automated analysis components in this section. The user (sonographer) would acquire a short video sweep of the fetal head using an ultrasound probe. Such an acquisition is easy to acquire with minimal training. In Fig. 2, engine A extracts TC frames from the video, and engine B localizes the cerebellum and measures the TCD to enable estimation of the GA. Importantly, the user does not need to identify the TC plane or make a manual TCD measurement.

2.2.1 Automated image frame detection (engine A)

There are various ways to perform automatic fetal ultrasound standard plane detection reported in the literature, including methods based on random forests²² and convolutional neural networks (CNNs).^{21,23,25,26,32} Fetal TC plane detection accuracies of the order of 90% and higher have been reported in a number of papers.^{22,23} In this work, we implemented a standard CNN, based on a variation of the well-known AlexNet architecture³³ that poses TC plane detection as a binary classification problem. A schematic of the CNN architecture is shown in Fig. 3. The ultrasound image is replicated three times as input. The CNN architecture consists of four convolutional layers followed by a fully connected layer. Subsequently, a softmax layer is used to produce the class probabilities. Network weights are initialized with values trained on the ImageNet dataset.³⁴ Using a pretrained network for initialization has been shown to significantly speed up training and improve classification accuracy of a network relative to choosing random values in other medical imaging applications³⁵ as well as fetal ultrasound image plane classification.²⁹

2.2.2 Automated image analysis module (engine B)

To perform GA estimation based on the TCD, an accurate diameter must be automatically measured. We formulate this as an image segmentation problem and train a fully convolutional neural network (FCN)³⁶ to segment the cerebellum, hereafter called the TC structure, from which the TCD can be estimated.

An FCN is a variant of a CNN where the last fully connected layer(s) are typically substituted by one or more convolution layers. In addition, transposed convolution layers are added to upsample the output of the convolutional layers and hence the output of a FCN is a segmentation mask that is the same size as the original image. In such an architecture, the entire network can be trained end-to-end. Details of the FCN architecture design and parameters employed are given in Fig. 4. A softmax loss function was used. The output of the pool3 layer is summed with the 2× upsampled prediction layer, before a transposed convolution layer with stride 8 is applied on the resulting feature maps to obtain the segmentation map. In the experiments reported in Sec. 3.2, we compare FCN architectures with different input image sizes and with and without transfer learning initialization.

Fig. 1 Schematic and representative ultrasound image for (a) the TV standard imaging plane and (b) the TC standard imaging plane.

Fig. 2 Analysis framework schematic. A short ultrasound video is acquired using a low-cost ultrasound device. Next, TC plane frames are extracted from the ultrasound video in engine A. Engine B then localizes the TC structure and performs automated TCD estimation from which the GA can be estimated via an equation.

Fig. 3 TC plane classification method schematic. An ultrasound frame fed to the CNN architecture is classified as a TC plane or otherwise.

The training set used to derive the FCN segmentation model consisted of 500 images from dataset B. This is a relatively small dataset for training a deep learning model. Data augmentation was employed to boost segmentation performance. This included random translation by 25% of image dimensions, random zooming by up to 20%, random rotations between -10 deg and 10 deg, and changes of image intensity to mimic random brightness and contrast variations.

A typical input image, manually segmented TC structure (ground truth), and automatically segmented TC structure mask are shown in Fig. 5.

The TCD can be automatically measured from the automatic segmentation mask. This is achieved by finding the major axes of the segmented cerebellum structure, assuming it is elliptical in shape. Although this is an approximation, experimentally this

Fig. 4 Schematic representation of the TC structure segmentation FCN-based network architecture.

Fig. 5 (a) A fetal ultrasound image where the TC structure is visible, (b) the ground truth segmentation label for the TC structure overlaid on the image, and (c) the automated TC structure segmentation result.

assumption turns out to lead to good TCD estimates compared with manual measurement (see the results in Sec. 3). Given the size of a pixel in mm, the location of the two pixels at the extremes of the major axis is transformed from pixel space to physical space. The Euclidean distance is then computed between these two points to estimate the TCD in mm.

GA can be calculated from the TCD estimate in millimeters using established parametric equations. In this work, we used the equation from the INTERGROWTH-21st study:³⁷

$$GA = 4.777048 + (0.724321 * TCD). \quad (1)$$

Here, GA is the gestational age in weeks and TCD is the transcerebellar diameter in millimeters.

3 Experimental Evaluation

Experiments were designed to evaluate the accuracy of: (1) TC plane detection, (2) TC structure segmentation for different FCN variants, and (3) automatic versus manual TCD-based GA estimation. We also present a first result of applying the developed image analysis methods to low-cost ultrasound probe video.

3.1 TC Plane Detection

The TC plane classification algorithm was trained for 1000 iterations with 128 images in each batch, on a training dataset of 3000 images using an NVIDIA TitanXp GPU. The training dataset consisted of 1500 images of the TC plane as positive examples and 1500 images of the TV plane as negative examples. The resulting model had around 6.4 million parameters. Subsequently, the validation set consisted of 1000 unseen data samples with 500 positive and 500 negative images. On the validation set, the algorithm precision and recall for TC plane detection were both 0.99, and the average accuracy 99%. Figure 6 shows the first 500 iterations of the training loss and of the validation accuracy.

3.2 Transcerebellum Segmentation/Detection

The accuracy of the TC structure segmentation algorithm was evaluated on, and compared with, images of varying size and FCN architectures with and without transfer learning to see if this made a difference. Where employed, transfer learning was used to initialize the first two layers of the network prior to training. It is important to note that the cerebellum is a small structure in the fetal brain. This translates to a high imbalance between the pixels that correspond to the class of interest (the TC structure) versus the background pixels (Fig. 5). Therefore, considering

solely the pixel segmentation accuracy would not be a suitable metric. Instead, we utilize the metrics introduced in Ref. 33, namely the mean accuracy and the mean intersection-over-union (IOU), that takes class imbalance into account. In essence, the IOU metrics are preferred since they control for the total number of pixels for a class. Referring to Table 1, overall accuracy measures the correctly predicted pixels of the foreground class, class c (the cerebellum), divided by the total number of pixels. Mean accuracy measures the correctly predicted pixels of class c divided by the total number of pixels of that class. Mean IOU is the IOU averaged over all classes.

Table 1 presents TC structure segmentation results for input images of varied size where the frames are downsampled from 128×192 pixels to 96×128 pixels to 64×96 pixels. In addition, in each case, the effect of transfer learning (using a pretrained network) is shown.

As can be seen from Table 1, the highest mean accuracy is achieved for an input size of $128 \times 192 \times 3$ pixels. Furthermore, as the size of the input image decreases to 96×128 and 64×96 , the mean accuracy achieved is also decreased to 85.07% and 80.86%, respectively. In each case, and for both metrics, transfer learning had a small positive impact in boosting the accuracy. The best segmentation task model had around 38 million parameters.

Fig. 6 TC plane classification training loss and validation accuracy for the first 500 iterations.

Table 1 Cerebellum (transcerebellum structure) segmentation accuracy. The input to the network is height $\times$ width $\times$ 3. Overall accuracy measures the correctly predicted pixels of class c (the cerebellum) divided by the total number of pixels, mean accuracy measures the correctly predicted pixels of class c divided by the total number of pixels of that class, and mean intersection over union is the intersection over union averaged over all classes. TF, TransFer learning.

Frame size (height $\times$ width)	Overall accuracy (%)	Mean accuracy (%)	Mean intersection over union (%)
(64 $\times$ 96 $\times$ 3)NoTF	96.78	79.59	73.25
(64 $\times$ 96 $\times$ 3)WithTF	96.95	80.86	77.57
(96 $\times$ 128 $\times$ 3)NoTF	97.38	83.62	77.66
(96 $\times$ 128 $\times$ 3)WithTF	97.52	85.07	78.93
(128 $\times$ 192 $\times$ 3)NoTF	97.82	86.30	81.00
(128 $\times$ 192 $\times$ 3)WithTF	97.89	86.99	81.62

3.3 TCD Measurement and GA Estimation

Automatically derived TCD estimates were compared with manual measurements (the ground truth) for dataset B ($n = 3236$ images).

Segmentation failure associated with low image quality will lead to an invalid measurement. Based on visual inspection, a measurement difference threshold of 5 mm was empirically selected to capture this source of error. Using this value, the algorithm failed in 35 cases only (1% of the total cases). For the full dataset, the mean of the manual measurements was 20.6 mm with a standard deviation (stdev) of 1.2 mm, compared with the mean of the automated measurement of 20.3 mm (stdev 2.9 mm) (p -value < 0.01). When excluding the 35 failure cases, the mean manual measurements was 20.6 mm (stdev 1.2 mm) while the mean of the automated measurement was 20.3 mm (stdev 1.6 mm) (p -value < 0.01), suggesting automated measurement accuracy is very good and that the algorithm is stable on the majority of cases.

Equation (1) was used to convert the TCD estimates in millimeters to GA in weeks. This allowed any error to be expressed in terms of GA units (weeks). The mean manual GA estimation was found to be 19.7 weeks (stdev 0.9) compared with automated estimation of 19.5 weeks (stdev 2.1) (p -value < 0.1) and, when excluding the 35 failure cases, 19.5 weeks (stdev 1.2) (p -value < 0.1). Figure 7 shows how manual and automated GA estimates are correlated ($n = 3201$ estimates excluding the 35 failure cases). Figure 8 shows a Bland–Altman plot to visualize the agreement between the manual and automated GA estimation methods. Again, we conclude that there is a good agreement between the two estimates.

3.4 Application of Methods to a Video from a Low-Cost Ultrasound Probe

In this section, we report a first result of applying the machine-learning-based GA estimation methods to a single ultrasound video from a low-cost ultrasound probe. Note that the algorithms were applied “as is” with no attempt to refine the deep learning models to the local properties of the low-cost ultrasound video.

Fig. 7 Manual versus automatic GA estimates in weeks of GA ($n = 3201$).

Fig. 8 Bland–Altman plot to show the agreement between the manual and automatic GA estimation ($n = 3201$).

For this experiment, we used the Konted Gen1 C10R (3.5-MHz convex) wireless ultrasound probe for ultrasound video acquisition (Konted, Beijing, China). A single ultrasound video was acquired following the acquisition protocol advocated for use with the machine-learned methods, namely the sonographer found the TV plane and then slowly moved the probe from the TV to include the TC over ~ 8 s. Retrospectively, an experienced sonographer viewed the video frame-by-frame and located the frame they considered to be the best TC plane based on clinical guidelines for TC standard plane definition.³⁸ A sonographer manually measured the TCD in this frame. For comparison, we also had access to the clinical hospital scan TCD measurement on the same subject made using a high-end ultrasound machine (GE Voluson E8).

Fig. 9 A low-cost ultrasound probe TC plane image: (a) original image and (b) image with automated segmentation overlaid—the TC structure is undersegmented.

The TC plane detection algorithm reported in this paper was applied to all video frames. The TC structure segmentation and TCD estimation algorithms reported in this paper were applied to the frame identified by the sonographer as the best TC plane.

The TC plane detection algorithm found the correct TC plane with a high class probability, but TC plane detection had a high false-positive rate overall. TCD automated measurement (26.2 mm) was an underestimation compared with manual measurement (34.9 mm) and the hospital measurement (36.2 mm). The underestimation was found to be due to undersegmentation of the TC structure, see Fig. 9.

4 Discussion

We have presented a pipeline of the key image analysis tasks suitable for a point-of-care low-cost ultrasound-based fetal GA solution. The solution design is intended to be (1) usable by a minimally trained, LMIC healthcare professional using a simple scanning protocol and (2) suitable for a wide range of GA by using the TCD as the basis for GA estimation.

Although the primary objective is to provide a point-of-care GA estimate, the proposed GA estimation device is designed to enable the user to separate data acquisition from the automated analysis stage, if desired. This could be helpful if either multiple scans are taken or multiple subjects are scanned, as analysis can be carried out at a later time.

The deep learning architectures were chosen as they are readily available CNN architectures for proof-of-concept demonstration of the complete pipeline for this application. However, they are large CNN models (millions of parameters) and are not well suited for deployment on smartphones or tablets with modest memory. More recent state-of-the-art deep learning architectures that result in compact models, such as MobileNet³⁹ and ShuffleNet,⁴⁰ will be investigated in future work but the trade off between accuracy and model size will need to be understood.

We have proposed a simple realization of the tool although thus far the analysis components have only been thoroughly tested separately as a proof of principle due to the current lack of availability of large volumes of low-cost ultrasound probe data. However, in Sec. 3.3, we reported on the application of the developed models to one of the early ultrasound videos we have acquired using a low-cost probe. These results are encouraging and show that our proposed automated methodology works in principle on images and video from this low-cost probe, but, perhaps not unexpectedly, the machine learning models will require fine tuning on video from this domain to achieve clinically accepted results. For the reported result, we chose a video for which the TC was visible to the human eye

in the video. The image appearance of the low-cost probe ultrasound video was different from the higher end ultrasound probe, but the sonographer was confident that they could see the TC structure in the video. In general, the visibility of the TC may be lower than a standard clinical probe, but research needs to be conducted to either confirm or refute this possible limitation.

More work is likely to be required on standard plane verification if the tool is to be used by healthcare professionals who do not have the skill to find an accurate standard plane. Attention or saliency maps may be useful here as suggested in Refs. 24 and 30.

The automated TCD measurement algorithm has been well validated on a large data set and achieved very good accuracy compared with manual measurement on the same images. We believe this to be the first time automated TCD estimation that has been demonstrated at scale.

A natural question to ask is whether the accuracy is good enough for clinical use. The clinically acceptable age difference (error) depends on clinical use of a GA estimation so there is no single answer to the question of what is clinically acceptable. Indeed, the improvement of GA estimation is an intense field of clinical research.

Therefore, it seems reasonable to say that an “acceptable” difference in GA estimation is what is acceptable compared with current standards—clinical measurement of the fetus. According to Ref. 41, these are ± 4 to 7 days at around 12 weeks, 10 to 12 days at around 20 weeks, ± 14 days at 28 weeks, and ± 18 to 21 days at 34 weeks. Our results, which are on conventional ultrasound images, are within these bounds for the second trimester, but it is yet to be proven whether the same accuracy can be achieved using a low-cost ultrasound probe and across a larger GA range. The contributing factors to accuracy in our particular case will be first the image plane finding (sonographer skill), second structure appearance (fetal and maternal factors), and third the accuracy of automated segmentation. The relative importance of these contributions will be investigated in future work when we have a working low-cost ultrasound-based device in the field.

Current work is integrating the components together into a single android-based device suitable for a minimally trained user, as well as there are plans to conduct a usability study of the next-generation solution in the field, which will be reported in a future publication.

5 Conclusion

The main contribution of this paper is to describe and implement key components of a computational framework for estimating the fetal GA from an ultrasound video sweep of the fetal brain

as the basis of a POCUS-based GA estimation solution for LMICs. Future work will implement the fully integrated device and conduct translational research usability studies within the context of the sub-Saharan African PRECISE Network.⁴²

Disclosures

The authors declare no conflicts of interest with the research described in the paper.

Acknowledgments

This research was funded by the University of Oxford EPSRC Impact Acceleration Award (EP/K503769/1), and a Bill & Melinda Gates Foundation Grand Challenges Exploration award (Project No. 1128591). MAM acknowledges the RCUK Centre for Doctoral Training in Healthcare Innovation for a doctoral research studentship (EP/G036861/1). JAN acknowledges EPSRC Grants EP/M013774/1, EP/L505316/1, EP/R013853/1, and ERC Advanced Grant No. 694581 (PULSE). Aris T. Papageorghiou was supported by the National Institute for Health Research (NIHR) Oxford Biomedical Research Centre (BRC). The low-cost ultrasound video was acquired as part of the MRC GCRF Programme PRECISE Network Grant (MR/P027938/1). Statement on Use of Data from Human Subjects: The authors confirm that the appropriate institutional approval was received to use the clinical images used in the research described in this paper.

References

- World Health Organization, "Fact sheets: maternal mortality," 2018, <https://www.who.int/news-room/fact-sheets/detail/maternal-mortality>.
- R. Hershkovitz, E. Sheiner, and M. Mazor, "Ultrasound in obstetrics: a review of safety," *Eur. J. Obstet. Gynecol. Reprod. Biol.* **101**(1), 15–18 (2002).
- R. Anderson et al., "The midwife's assistant: designing integrated learning tools to scaffold ultrasound practice," in *Proc. Fifth Int. Conf. Inf. and Commun. Technol. and Dev.*, pp. 200–210 (2012).
- W. Brunette et al., "Portable antenatal ultrasound platform for village midwives," in *Proc. First ACM Symp. Comput. for Dev.*, pp. 23:1–23:10 (2010).
- T. L. A. van den Heuvel et al., "A step towards measuring the fetal head circumference with the use of obstetric ultrasound in a low resource setting," *Proc. SPIE* **10139**, 101390V (2017).
- K. K. DeStigter et al., "Low-cost teleradiology for rural ultrasound," in *IEEE Global Humanitarian Technol. Conf.*, pp. 290–295 (2011).
- H. Blencowe et al., "National, regional, and worldwide estimates of preterm birth rates in the year 2010 with time trends since 1990 for selected countries: a systematic analysis and implications," *Lancet* **379**(9832), 2162–2172 (2012).
- M. de Onis, M. Blössner, and J. Villar, "Level and patterns of intrauterine growth retardation in developing countries," *Eur. J. Clin. Nutr.* **52**(Suppl 1), S5–S15 (1998).
- S. Campbell et al., "Routine ultrasound screening for the prediction of gestational age," *Obstet. Gynecol.* **65**(5), 613–620 (1985).
- T. H. Nguyen et al., "Evaluation of ultrasound-estimated date of delivery in 17 450 spontaneous singleton births: do we need to modify Naegele's rule?" *Ultrasound Obstet. Gynecol.* **14**(1), 23–28 (1999).
- R. Napolitano et al., "Pregnancy dating by fetal crown-rump length: a systematic review of charts," *BJOG* **121**(5), 556–565 (2014).
- M. R. Chavez et al., "Fetal transcerebellar diameter measurement for prediction of gestational age at the extremes of fetal growth," *J. Ultrasound Med.* **26**(9), 1167–1171 (2007).
- E. T. Kim et al., "Obstetric ultrasound use in low and middle income countries: a narrative review," *Reprod. Health* **15**(1), 129 (2018).
- I. Sarris et al., "Standardization of fetal ultrasound biometry measurements: improving the quality and consistency of measurements," *Ultrasound Obstet. Gynecol.* **38**(6), 681–687 (2011).
- S. Z. Wanyonyi et al., "Image-scoring system for crown–rump length measurement," *Ultrasound Obstet. Gynecol.* **44**(6), 649–654 (2014).
- M. J. Rijken et al., "Obstetric ultrasound scanning by local health workers in a refugee camp on the Thai–Burmese border," *Ultrasound Obstet. Gynecol.* **34**(4), 395–403 (2009).
- T. L. A. van den Heuvel et al., "Comparison study of low-cost ultrasound devices for estimation of gestational age in resource-limited countries," *Ultrasound Med. Biol.* **44**(11), 2250–2260 (2018).
- ISUOG, "ISUOG Education Committee recommendations for basic training in obstetric and gynecological ultrasound," *Ultrasound Obstet. Gynecol.* **43**(1), 113–116 (2014).
- A. I. L. Namburete et al., "Learning-based prediction of gestational age from ultrasound images of the fetal brain," *Med. Image Anal.* **21**(1), 72–86 (2015).
- S. Rueda et al., "Evaluation and comparison of current fetal ultrasound image segmentation methods for biometric measurements: a grand challenge," *IEEE Trans. Med. Imaging* **33**(4), 797–813 (2014).
- H. Chen et al., "Standard plane localization in fetal ultrasound via domain transferred deep neural networks," *IEEE J. Biomed. Health Inf.* **19**(5), 1627–1636 (2015).
- M. Yaqub et al., "Guided random forests for identification of key fetal anatomy and image categorization in ultrasound scans," *Lect. Notes Comput. Sci.* **9351**, 687–694 (2015).
- C. F. Baumgartner et al., "SonoNet: real-time detection and localisation of fetal standard scan planes in freehand ultrasound," *IEEE Trans. Med. Imaging* **36**(11), 2204–2215 (2017).
- J. Schlemper et al., "Attention-gated networks for improving ultrasound scan plane detection," arXiv:1804.05338 (2018).
- H. Chen et al., "Automatic fetal ultrasound standard plane detection using knowledge transferred recurrent neural networks," *Lect. Notes Comput. Sci.* **9349**, 507–514 (2015).
- H. Chen et al., "Ultrasound standard plane detection using a composite neural network framework," *IEEE Trans. Cybern.* **47**(6), 1576–1586 (2017).
- C. P. Bridge, C. Ioannou, and J. A. Noble, "Automated annotation and quantitative description of ultrasound videos of the fetal heart," *Med. Image Anal.* **36**, 147–161 (2017).
- M. A. Maraci et al., "A framework for analysis of linear ultrasound videos to detect fetal presentation and heartbeat," *Med. Image Anal.* **37**, 22–36 (2017).
- Y. Gao, M. A. Maraci, and J. A. Noble, "Describing ultrasound video content using deep convolutional neural networks," in *IEEE 13th Int. Symp. Biomed. Imaging (ISBI)*, pp. 787–790 (2016).
- Y. Gao and J. A. Noble, "Detection and characterization of the fetal heartbeat in free-hand ultrasound sweeps with weakly-supervised two-streams convolutional networks," *Lect. Notes Comput. Sci.* **10343**, 305–313 (2017).
- T. L. A. van den Heuvel et al., "Combining automated image analysis with obstetric sweeps for prenatal ultrasound imaging in developing countries, imaging for patient-customized simulations and systems for point-of-care ultrasound," *Lect. Notes Comput. Sci.* **10549**, 105–112 (2017).
- M. Yaqub et al., "A deep learning solution for automatic fetal neurosonographic diagnostic plane verification using clinical standard constraints," *Ultrasound Med. Biol.* **43**(12), 2925–2933 (2017).
- A. Krizhevsky, I. Sutskever, and G. A. Hinton, "ImageNet classification with deep convolutional neural networks," *Adv. Neural Inf. Process. Syst.* **25**, 1097–1105 (2012).
- O. Russakovsky et al., "ImageNet large scale visual recognition challenge," *Int. J. Comput. Vision* **115**(3), 211–252 (2015).
- H.-C. Shin et al., "Deep convolutional neural networks for computer-aided detection: CNN architectures, dataset characteristics and transfer learning," *IEEE Trans. Med. Imaging* **35**(5), 1285–1298 (2016).
- J. Long et al., "Fully convolutional networks for semantic segmentation," in *Proc. IEEE Conf. Computer Vision Pattern Recognition*, pp. 3431–3440 (2015).
- V. Donadono et al., "Development of normal fetal brain structures: standards based on ultrasound measurements from the fetal growth longitudinal study of the INTERGROWTH-21st Project," *Ultrasound Obstet. Gynecol.* **50**(S1), 10–11 (2017).
- D. Paladini et al., "Sonographic examination of the fetal central nervous system: guidelines for performing the 'basic examination' and the 'fetal neurosonogram,'" *Ultrasound Obstet. Gynecol.* **29**(1), 109–116 (2007).

39. A. G. Howard et al., "MobileNets: efficient convolutional neural networks for mobile vision applications," arXiv:1704.04861 (2017).
40. X. Zhang et al., "ShuffleNet: an extremely efficient convolutional neural network for mobile devices," arXiv:1707.01083 (2017).
41. A. T. Papageorghiou et al., "Ultrasound-based gestational-age estimation in late pregnancy," *Ultrasound Obstet. Gynecol.* **48**(6), 719–726 (2016).
42. www.precisenetwork.org

Mohammad A. Maraci received his MSc degree in databases and web-based systems and subsequently completed his PhD in health-care innovation at the Institute of Biomedical Engineering (IBME), University of Oxford, United Kingdom, on the topic of ultrasound image analysis methods for resource-constrained environments. He maintained an active interest in affordable healthcare activities while a PhD student and subsequently held postdoctoral research fellow positions at the IBME considering translation of obstetric ultrasound images analysis methods into LMIC settings.

Mohammad Yaqub received his PhD in machine-learning-based fetal ultrasound image analysis from the IBME, University of Oxford, United Kingdom, and then worked as a postdoctoral research fellow at the IBME on cardiac and fetal ultrasound image analysis. He left in 2016 to work at Intelligent Ultrasound Ltd., where he is currently a vice-president of engineering. He has maintained an academic interest in teaching and currently holds an academic visitor appointment at Oxford University.

Rachel Craik received her first degree in biology from York University and holds research coordinator posts at King's College London and Oxford University. She coordinated the Oxford-led INTERBIO-21st Study, a global initiative investigating the growth and development of infants from conception to the age of 2 years in high-risk settings. Currently, she co-ordinates the KCL-led PRECISE study in sub-Saharan Africa and a number of Oxford-led global ultrasound clinical and translational projects including this one.

Sridevi Beriwal is a clinical research fellow in the Nuffield Department of Women's and Reproductive Health at the University of

Oxford. She works on the PRECISE global health ultrasound project with IBME engineers and clinical teams in the Gambia, Kenya, and Mozambique.

Alice Self is a clinical research fellow in the Nuffield Department of Women's and Reproductive Health at the University of Oxford with a special interest in global health. She primarily works on the CALOPUS project; an interdisciplinary collaboration between NDWRH, IBME and THSTI, New Delhi, aiming to develop more affordable, point-of-care, automated ultrasound technology for use in LMICs, and is an Oxford clinical team member on the PRECISE study.

Peter von Dadelszen, an obstetrician (maternal-fetal medicine), is a Professor of Global Women's Health, King's College London. He has trained, certified, and practiced in New Zealand, Canada, and the United Kingdom. With Professor Laura Magee, he leads and collaborates on research into pregnancy hypertension, fetal growth, and stillbirth, from basic science to clinical epidemiology and health services research to global health, primarily through the Gates Foundation-funded PRE-EMPT and UKRI-funded PRECISE initiatives.

Aris Papageorghiou is a clinician scientist in fetal medicine based in Oxford and London. His research is in the areas of maternal, fetal, and perinatal health using diverse methods—basic science, epidemiology, trials, and implementation science. This has included large global projects that aim to improve phenotyping and assess fetal/infant growth and development. His research interest is improving pregnancy care in low-resource settings with AI-assisted ultrasound imaging.

J. Alison Noble is the Oxford University Technikos Professor of Biomedical Engineering. Her research group is best known for its interdisciplinary research at the interfaces of ultrasound imaging, computational (machine-learning-based) image analysis, and clinical medicine. She has published widely on the topic and holds a European Research Council Advanced Grant on next-generation ultrasound systems. She is a former president of the MICCAI Society and a fellow of the UK Royal Society.