

Automatic Probe Movement Guidance for Freehand Obstetric Ultrasound

Richard Droste¹, Lior Drukker², Aris T. Papageorghiou², and J. Alison Noble¹

¹ Institute of Biomedical Engineering, University of Oxford, Oxford, UK
`richard.droste@eng.ox.ac.uk`

² Nuffield Department of Women's & Reproductive Health,
University of Oxford, Oxford, UK

Abstract. We present the first system that provides real-time probe movement guidance for acquiring standard planes in routine freehand obstetric ultrasound scanning. Such a system can contribute to the world-wide deployment of obstetric ultrasound scanning by lowering the required level of operator expertise. The system employs an artificial neural network that receives the ultrasound video signal and the motion signal of an inertial measurement unit (IMU) that is attached to the probe, and predicts a guidance signal. The network termed US-GuideNet predicts either the movement towards the standard plane position (goal prediction), or the next movement that an expert sonographer would perform (action prediction). While existing models for other ultrasound applications are trained with simulations or phantoms, we train our model with real-world ultrasound video and probe motion data from 464 routine clinical scans by 17 accredited sonographers. Evaluations for 3 standard plane types show that the model provides a useful guidance signal with an accuracy of 88.8% for goal prediction and 90.9% for action prediction.

Keywords: Fetal ultrasound · Probe guidance · Ultrasound navigation

1 Introduction

Ultrasound scanning is an indispensable diagnostic tool in obstetrics due to its safety, real-time results and low cost. At the same time, many women in developing countries do not receive a single ultrasound examination throughout their pregnancy due to a lack of skilled operators [20]. The main tasks of ultrasound scanning are the acquisition, examination/verification and interpretation of pre-defined standard anatomical planes that enable the detection of fetal abnormalities. Systems that provide assistance for or automate these tasks have the potential of enabling worldwide access to ultrasound scanning by reducing the level of necessary expertise. Standard plane examination/verification and interpretation are largely standardized [18], can be performed remotely [3], and can be facilitated through automated image analysis [24]. Freehand standard plane acquisition, on the other hand, is harder to facilitate/automate since it is not standardized and requires interaction with the mother. It demands years of training and is the rate-limiting step even for experienced sonographers [1].




Fig. 1. System overview. *Left:* An operator performs ultrasound scanning with a routine clinical setup while the motion of the probe is recorded with an IMU. *Bottom right:* The *US-GuideNet* receives the IMU motion signal and the ultrasound video signal as inputs and outputs a real-time probe movement guidance signal. *Top right:* Attachment of the IMU to the ultrasound probe and IMU coordinate system.

To address this issue, we present the first system that provides real-time probe movement guidance for fetal standard plane acquisition in routine freehand obstetric ultrasound scanning. An overview of the system is presented in Fig. 1. An artificial neural network termed *US-GuideNet* receives the ultrasound video signal alongside the signal of a motion sensor that is attached to the ultrasound probe, and outputs probe movement guidance that directs the operator towards the desired standard plane. No specialized equipment is required: The motion sensor is a common inertial measurement unit (IMU) that is attached to the probe of a standard clinical ultrasound machine. Further, the *US-GuideNet* neural network is designed to be extremely lightweight and can run real time inference on a CPU. Behavioral cloning (BC), a type of imitation learning, has emerged as a powerful technique to train neural networks to perform complex real-world tasks such as autonomous driving [14]. Here, we collect 5079 demonstrations of standard plane acquisitions from 464 2nd- and 3rd-trimester scans acquired by 17 accredited sonographers and implement two settings of BC: 1) For *goal prediction*, the network predicts the movement that leads directly to the estimated position of the standard plane. 2) For *action prediction*, the network predicts the next movement that the expert would perform.

2 Related Work

Various approaches have been proposed to address the difficulty of ultrasound standard plane acquisition.

Robotic Ultrasound. Human-controlled robotic systems have been developed that allow experienced sonographers to perform obstetric ultrasound exams remotely [22]. Automated robotic systems have been proposed for highly structured tasks such as finding planes of motionless objects [11,8] or the human liver [13]. However, despite on-going efforts [23], no robotic system has been proposed that can automate the complex task of obstetric ultrasound scanning.

Simplified Acquisition Protocols. Instead of assisting operators to acquire typical freehand 2D standard planes, previous work has proposed to automatically extract standard planes from data that are acquired with a simplified protocol, such as 3D ultrasound volumes [17,7] or linear sweeps over the maternal abdomen [10]. Moreover, IMUs have been used to acquire 3D ultrasound volumes with 2D probes [5,16]. However, these methods are applicable only for a subset of standard planes (fetal abdomen and head) and the standard plane quality is not up to par with typical freehand scanning.

Phantoms and Simulated Environments. Recent studies have proposed learning based systems that are trained to acquire ultrasound planes in simplified environments. One study proposes an algorithm that learns to find a view of the adult heart in a grid of pre-acquired ultrasound images [12]. Moreover, learning-based systems have been proposed in which a robotic actuator finds predefined views of simple tissue phantoms [6] or a fetal US phantom [21]. However, a fetus in the mother’s womb is a dynamic and highly variant object that can not be well-represented with static simulations or a phantom. Furthermore, these algorithms are purely image-based and therefore rely on the exact execution of the predicted actions, which is only possible within a simulation or with a robotic system. Here, we train a guidance algorithm with video and probe motion data from a large number of real-world expert demonstrations from routine scanning. Moreover, our algorithm receives the real-time probe motion signal and can therefore react to the movements of a human operator.

3 Method

Fig. 1 presents an overview of the proposed system. An operator performs routine obstetric ultrasound scanning with a standard clinical machine. An inertial measurement unit (IMU) motion sensor is attached to the ultrasound probe and an on-board attitude and heading reference system (AHRS) estimates the sensor’s orientation in the earth coordinate system. The motion sensor signal and the machine video signal are input into a neural network, *US-GuideNet*, that outputs a 3D rotation of the probe that guides the operator towards the standard plane. The network training method is described in Section 3.1, the network architecture is detailed in Section 3.2 and implementation details are provided in Section 3.3.

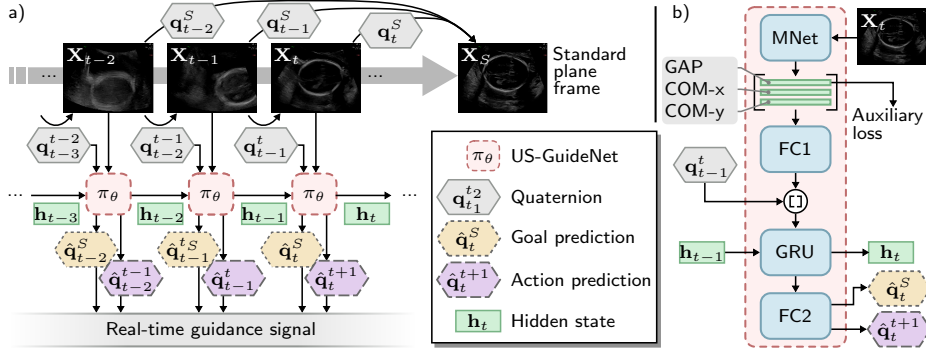


Fig. 2. a) Proposed behavioral cloning framework. b) *US-GuideNet* architecture.

3.1 Learning from Expert Demonstrations

We pose the problem of training a neural network to predict a probe guidance signal as a behavioral cloning problem. That is, we record standard plane acquisition demonstrations from several experts for a large number of patients and train the network to replicate the demonstrated behavior. In general, a standard plane acquisition consists of live B-mode scanning followed by *freezing* the ultrasound video and optionally selecting a previous frame with the desired appearance from a *cine-buffer*. We define the finally selected frame as the standard plane.

Problem Formulation. Fig. 2 a) presents the formulation of the learning problem. Let $\{\mathbf{X}_t \in \mathbb{R}^{H \times W} | t \in \mathcal{T}\}$ be ultrasound video frames of a standard plane acquisition, temporally downsampled to 6 Hz, with resolution $H \times W$ and frame indices $\mathcal{T} = \{i\}_{i=0}^F$, where F is the *freeze* frame index. Moreover, let \mathbf{X}_S be the standard plane with index $S \in \mathcal{T} \setminus \{0\}$. Finally, let $\mathbf{q}_t = [q_w, q_x, q_y, q_z]^\top$ be the probe orientation quaternion of frame \mathbf{X}_t and $\mathbf{q}_t^{t_2} := \mathbf{q}_t^* \mathbf{q}_{t_2}$ the probe rotation quaternion from frame \mathbf{X}_t to frame \mathbf{X}_{t_2} , where \mathbf{q}^* is the conjugate. We represent orientations with quaternions since they can be smoothly interpolated without discontinuities or singularities, and are numerically stable and computationally efficient [15]. Euler angles, in contrast, another popular representation of rotations, suffer from discontinuities such as the *gimbal lock*. We do not consider probe translation in this work since the IMU is not suitable to estimate it accurately.

Behavioral Cloning. We train a policy network $\pi_\theta : s_t \mapsto u_t$ termed *US-GuideNet* with parameters θ that maps the state s_t at time step $t \in \mathcal{T}$ to an action u_t . We define the state as the tuple $s_t := (\mathbf{X}_t, \mathbf{q}_{t-1}^t, \mathbf{h}_{t-1})$, where \mathbf{h}_{t-1} is the hidden state of a recurrent neural network within π_θ . We explore the two different settings for the action u_t : *goal prediction* and *action prediction*. For *goal prediction*, the policy $\pi_\theta^g : s^t \mapsto \hat{\mathbf{q}}_t^S$ estimates the rotation from the current orientation to the orientation of the standard plane. If the estimated standard plane orientation is accurate, this policy is optimal, *i.e.*, it guides the operator directly to the standard plane. However, it is not guaranteed that enough

information has been seen at time t for an accurate estimation of the standard plane orientation. Therefore, we explore a second setting, *action prediction*, where the policy $\pi_\theta^a : s_t \mapsto \hat{\mathbf{q}}_t^{t+1}$ estimates the next rotation that the operator would perform. This policy aims to closely mimic the expert sonographer behavior.

Loss Function. During training, a demonstration is constructed from a subset of indices $\mathcal{T}_D \subset \mathcal{T}$ with start and end indices t_0 and T . Let $\hat{\mathbf{Q}} := [\hat{\mathbf{q}}_t^{t_2}]_{t=t_0}^T$ and $\mathbf{Q} := [\mathbf{q}_t^{t_2}]_{t=t_0}^T$ be the predicted and ground truth rotation quaternion sequences, with $t_2 \in \{S, t+1\}$ for *goal prediction* and *action prediction* respectively. We add an auxiliary output after the MNet in order to facilitate and regularize its training: Since we want the MNet to recognize the appearance of standard planes, we input the average pooled MNet features into a softmax layer that predicts the class probabilities of the SonoNet standard plane classifier [2] for each frame. Let $\hat{\mathbf{P}} = [\hat{\mathbf{p}}_t]_{t=t_0}^T$ be the auxiliary softmax output and $\mathbf{Y} = [\mathbf{y}_t]_{t=t_0}^T$ the SonoNet class probabilities, with $\hat{\mathbf{p}}_t, \mathbf{y}_t \in \mathbb{R}_{\geq 0}^{14}$. The total training loss \mathcal{L} is

$$\mathcal{L} = \sum_{t=t_0+W}^T \left\{ \underbrace{-\frac{1}{\|\hat{\mathbf{q}}_t^{t_2}\|} \hat{\mathbf{q}}_t^{t_2} \cdot \mathbf{q}_t^{t_2}}_{\text{Similarity loss}} + \underbrace{\alpha (1 - \|\hat{\mathbf{q}}_t^{t_2}\|^2)^2}_{\text{Norm loss}} \right\} - \beta \sum_{t=t_0}^T \underbrace{\mathbf{y}_t^\top \text{diag}(\mathbf{w}) \log(\mathbf{p}_t)}_{\text{Auxiliary loss}}$$

where \cdot denotes the dot-product, $\alpha, \beta \in \mathbb{R}_{>0}$ are scalar weighting factors, $\mathbf{w} \in \mathbb{R}_{\geq 0}^{14}$ is a weight vector that balances the SonoNet class probabilities, and W is a warm up time for the rotation prediction.

3.2 US-GuideNet Architecture

The *US-GuideNet* policy network receives the ultrasound video and probe motion signals and outputs predicted expert probe rotations as described in Section 3.1. We design the architecture for small time and space computational complexity (runtime and model size) such that it can run real-time inference on the CPU of an inexpensive computer. The network architecture is illustrated in Fig. 2 b). At each time step t , the ultrasound video frame X_t is fed into a MobileNet V2 (MNet) convolutional neural network [19], which consists of a cascade of lightweight depthwise-separable and pointwise convolutions. We use MNet with a width-multiplier of 0.5, *i.e.*, 50% reduced number of channels. Next, the dimensionality of the MNet output is reduced with a custom *ConcatPool* operation that preserves both semantic and spatial information by concatenating global average pooled (GAP) features with the x and y coordinates of the centers of mass (COM-x/y) of the feature maps. After reducing the features to 128 channels with a fully-connected layer *FC1*, they are concatenated with the current probe rotation quaternion \mathbf{q}_{t-1}^t and input into a gated recurrent unit [4] (GRU) with 132 input channels and 128 hidden channels. Finally *FC2*, a fully-connected layer with one 128-channel hidden layer, outputs the 4-dimensional probe rotation quaternion.

3.3 Experimental Setup

Data Acquisition. The data were acquired as part of the PULSE (Perception Ultrasound by Learning Sonographic Experience) project, a prospective study

of routine fetal ultrasound scans performed in all trimesters by accredited sonographers and fetal medicine doctors at the maternity ultrasound unit, Oxford University Hospitals NHS Foundation Trust, Oxfordshire, United Kingdom. The exams were performed on a GE Voluson E8 scanner (General Electric, USA) while the video signal of the machine monitor was recorded lossless at 30 Hz. The motion of each of two curved linear array transducer (2D) probes was recorded with a NGIMU IMU/AHRS (x-io Technologies Ltd., UK). Each motion sensor was attached to the cable outlet of the probe with a custom 3D-printed mount as shown in Fig. 1. The probe orientation quaternions were sampled at 400 Hz. This study was approved by the UK Research Ethics Committee (Reference 18/WS/0051) and written informed consent was given by all participating pregnant women and operators. In this paper we use ultrasound video and corresponding gaze data of 464 second and third trimester scans acquired by 17 accredited sonographers between May 2018 and February 2020.

Data Processing. We extracted the standard plane acquisitions from the ultrasound scans with a purpose-built program based on optical character recognition. For each of the 5079 resulting acquisitions, the program outputs the corresponding live B-Mode scanning segment, the *freeze* frame and the *cine-buffer*-corrected standard plane. In addition, the program labels acquisitions of the biometry standard planes: the femur standard plane (FSP), the abdominal standard plane (ASP) and the trans-ventricular plane (TVP) [18], which we use for evaluation. The acquisition duration was limited to 10 s before the standard plane. We automatically corrected any lag between the video and motion signals by correlating frame differences with probe motion, and manually verified the synchronization. The video frames were cropped such that the ultrasound machine’s graphical user interface was removed, and normalized to zero-mean and unit-variance. The scans are divided into 80% for training and 20% for testing.

Training. For each training epoch, a demonstration of 32 frames is randomly selected from each standard plane acquisition, which corresponds to a duration 5.3 s at 6 Hz. It is ensured that $\min_t \{\mathbf{q}_t \cdot \mathbf{q}_S\} \geq 0.7$ for each demonstration. The frames are augmented by randomly changing of the brightness, contrast and gamma by $\pm 10\%$ and randomly symmetrically cropping up to 20% of the frame border. The frames are then down-sampled to the network input resolution of 224×288 . The MNet is pre-trained via the auxiliary loss with a large number of ultrasound frames. The entire *US-GuideNet* neural network is then trained from the demonstrations for 20 epochs with the AdamW optimizer [9] with weight decay of 10^{-2} and initial learning rate of 0.001, which is decayed by a factor of 0.1 every 8 epochs. The batch size is set to 8 and the warm up time for the rotation loss to 1 s. After training with all demonstrations, the model is fine-tuned for each evaluation plane (FSP, ASP and TVP) separately for 16 epochs.

Evaluation and Baseline. We evaluate the trained model on the full-length standard plane acquisitions (clipped to 10 s before the standard plane). For each time step, we classify the predicted probe rotation as correct if and only if $\angle(\mathbf{q}_t \hat{\mathbf{q}}_t^{t_2}, \mathbf{q}_{t_2}) < \angle(\mathbf{q}_t, \mathbf{q}_{t_2})$, $t_2 \in \{S, t + 1\}$, *i.e.*, if applying the predicted rotation reduces the angle to the target orientation. As before, $t_2 = S$ for *goal prediction*

and $t_2 = t + 1$ for *action prediction*. As a baseline rotation prediction we use \mathbf{q}_{t-1}^t , *i.e.*, continuing in the current direction of rotation at each time step.

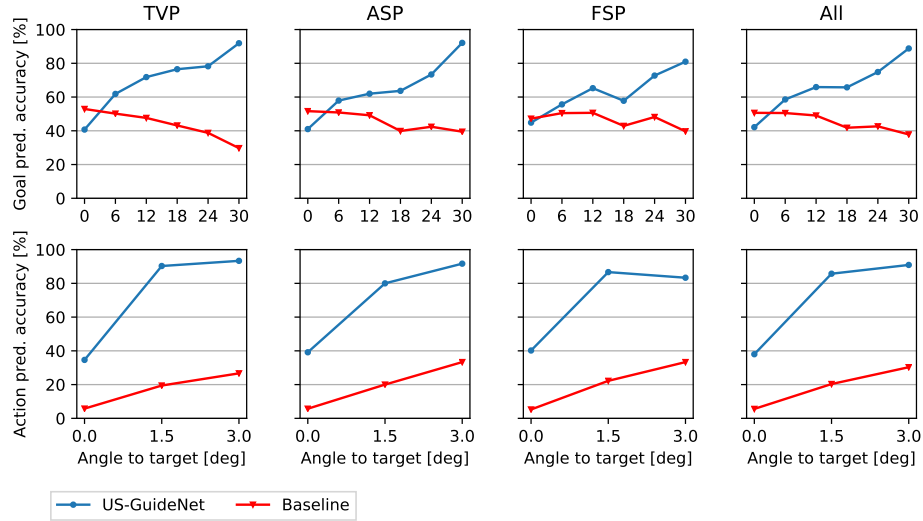


Fig. 3. Experimental results for the evaluated standard planes: TVP (head), ASP (abdomen) and FSP (femur). In addition, the overall accuracies are provided.

4 Results

The experimental results are shown in Fig. 3. The average accuracy of the guidance signal and baseline is evaluated for different ranges of the angular distance to the target (standard plane orientation for *goal prediction* or next probe position for *action prediction*). This enables the separation of the performance for coarse (large angular distance) and fine (low angular distance) adjustments. The x-axis of the individual plots provides the lower limits of the ranges, which extend to the next-higher x-axis value. Across the *action prediction* and *goal prediction* settings and all evaluated standard plane types, a common pattern can be observed that the accuracy of the guidance signal tends to increase with increasing angular distance to the standard plane.

Goal Prediction. The *goal prediction* accuracies are given in the upper row of Fig. 3. The guidance signal performs better than the baseline for any angle range $>6^\circ$. The accuracy of the guidance signal increases with increasing angular distances to the standard plane, ranging from 42.2% for angles 0° to 6° to 88.8% for angles $>30^\circ$, with 81.0% for the FSP and 92.1% for the ASP. The average baseline accuracy slightly declines towards higher angular distances.

Action Prediction. The guidance signal accuracy is higher than the baseline accuracy for all target distance ranges. The average guidance signal accuracy increases from 38.0% for angles 0° to 1.5° to 90.9% for distances $>3^\circ$. At angles $>3^\circ$, the largest accuracy is observed for the TVP with 93.3% and the lowest for the FSP with 83.3%. The average baseline accuracy slightly increases with increasing angular distances.

5 Discussion and Conclusion

The results presented in Section 4 demonstrate that the proposed probe guidance system for obstetric ultrasound scanning indeed provides a useful navigation signal towards the respective target, which is the standard plane orientation for *goal prediction* and the next expert movement for *action prediction*. The accuracy of *US-GuideNet* increases for larger differences to the target orientation, which shows that the algorithm is robust for guiding the operator towards the target orientation from distant starting points. For small distances, it is difficult to predict an accurate guidance signal since the exact target orientation may be subject to inter- and intra-sonographer variations or sensor uncertainty. The accuracy is similar for *goal prediction* and *action prediction* but slightly higher for action prediction at intermediate angles to the target, which can be explained by the fact that the action is always based on the previously seen frames, while the target position might be yet unknown.

The guidance signal accuracies are generally the highest for the abdominal and head standard planes (ASP and TVP). The accuracy for the femur standard plane (FSP) is slightly lower, which can be explained by the fact the the femur is part of an extremity and therefore subject to more fetal movement, which can make its final position unpredictable. Moreover, the FSP is defined via two anatomical landmarks—the distal and proximal ends of the femur—while the ASP and TVP are determined by the appearance of more anatomical structures [18]. This might make it more difficult for the model to predict the FSP position that was chosen by the operator, since it is subject to more degrees of freedom.

A limitation of our study is that we test our algorithm with pre-acquired data. However, in contrast with previous work [12,6,21] which uses simulations or phantoms, our proposed system is trained and evaluated on data from real-world routine ultrasound scanning. Moreover, instead of relying on the exact execution of the probe guidance as in previous work [12,6,21], our system reacts to the actual operator probe movements that are sensed with an IMU. This suggests that the system will perform well in future tests on volunteer subjects. In general, the accuracy of the predictions of *US-GuideNet* is evident from the large improvements over the baseline of simply continuing the current direction of rotation. While probe translation is not predicted due to IMU limitations, only the through-plane sweeping translation would usually change the view of the fetus while the sideways sliding and downwards/upwards translations would shift the fetal structure within the ultrasound image. In combination with the rotation guidance, this leaves one degree of freedom to be determined by the operator.

In conclusion, this paper presents the first probe movement guidance system for the acquisition of standard planes in freehand obstetric ultrasound scanning. Moreover, it is the first guidance system for any application of ultrasound standard plane acquisition that is trained with video and probe motion data from routine clinical scanning. Our experiments have shown that the proposed *US-GuideNet* network and behavioral cloning framework result in an accurate guidance system. These results will serve as a foundation for subsequent validation studies with novice operators. The proposed algorithm is lightweight which facilitates the deployment for existing ultrasound machines.

Acknowledgements. We acknowledge the ERC (ERC-ADG-2015 694581, project PULSE) and the NIHR Oxford Biomedical Research Centre.

References

1. Bahner, D.P., Blickendorf, J.M., Bockbrader, M., Adkins, E., Vira, A., Boulger, C., Panchal, A.R.: Language of Transducer Manipulation. *J. Ultrasound Med.* **35**(1), pp. 183–188 (2016)
2. Baumgartner, C.F., Kamnitsas, K., Matthew, J., Fletcher, T.P., Smith, S., Koch, L.M., Kainz, B., Rueckert, D.: SonoNet: Real-Time Detection and Localisation of Fetal Standard Scan Planes in Freehand Ultrasound. *IEEE Trans. Med. Imag.* **36**(11), pp. 2204–2215 (2017)
3. Britton, N., Miller, M.A., Safadi, S., Siegel, A., Levine, A.R., McCurdy, M.T.: Tele-Ultrasound in Resource-Limited Settings: A Systematic Review. *Front. Public Health* **7**, 244 (2019)
4. Cho, K., van Merriënboer, B., Bahdanau, D., Bengio, Y.: On the Properties of Neural Machine Translation: Encoder-Decoder Approaches. In: Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation (SSST-8), pp. 103–111 (2014)
5. Housden, R., Treece, G.M., Gee, A.H., Prager, R.W.: Calibration of an orientation sensor for freehand 3D ultrasound and its use in a hybrid acquisition system. *BioMed. Eng. OnLine* **7**(1), 5 (2008)
6. Jarosik, P., Lewandowski, M.: Automatic Ultrasound Guidance Based on Deep Reinforcement Learning. In: IEEE International Ultrasonics Symposium (IUS), pp. 475–478 (2019)
7. Li, Y., Khanal, B., Hou, B., Alansary, A., Cerrolaza, J.J., Sinclair, M., Matthew, J., Gupta, C., Knight, C., Kainz, B., Rueckert, D.: Standard Plane Detection in 3D Fetal Ultrasound Using an Iterative Transformation Network. In: International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI), pp. 392–400 (2018)
8. Liang, K., Rogers, A.J., Light, E.D., von Allmen, D., Smith, S.W.: Three-Dimensional Ultrasound Guidance of Autonomous Robotic Breast Biopsy: Feasibility Study. *Ultrasound Med. Biol.* **36**(1), pp. 173–177 (2010)
9. Loshchilov, I., Hutter, F.: Decoupled Weight Decay Regularization. In: International Conference on Learning Representations (ICLR) (2019)
10. Maraci, M.A., Bridge, C.P., Napolitano, R., Papageorghiou, A., Noble, J.A.: A framework for analysis of linear ultrasound videos to detect fetal presentation and heartbeat. *Med. Image Anal.* **37**, pp. 22–36 (2017)

11. Mebarki, R., Krupa, A., Chaumette, F.: 2-D Ultrasound Probe Complete Guidance by Visual Servoing Using Image Moments. *IEEE Trans. Robot.* **26**(2), pp. 296–306 (2010)
12. Milletari, F., Birodkar, V., Sofka, M.: Straight to the Point: Reinforcement Learning for User Guidance in Ultrasound. In: *Smart Ultrasound Imaging and Perinatal, Preterm and Paediatric Image Analysis (MICCAI Workshops)*, pp. 3–10 (2019)
13. Mustafa, A.S.B., Ishii, T., Matsunaga, Y., Nakadate, R., Ishii, H., Ogawa, K., Saito, A., Sugawara, M., Niki, K., Takanishi, A.: Development of robotic system for autonomous liver screening using ultrasound scanning device. In: *IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pp. 804–809 (2013)
14. Pan, Y., Cheng, C.A., Saigol, K., Lee, K., Yan, X., Theodorou, E., Boots, B.: Agile Autonomous Driving using End-to-End Deep Imitation Learning. In: *Robotics: Science and Systems (RSS)* (2019)
15. Pavlo, D., Feichtenhofer, C., Auli, M., Grangier, D.: Modeling Human Motion with Quaternion-based Neural Networks. *Int. J. Comput. Vis.* **128**, pp. 855–872 (2020)
16. Prevost, R., Salehi, M., Jagoda, S., Kumar, N., Sprung, J., Ladikos, A., Bauer, R., Zettinig, O., Wein, W.: 3D freehand ultrasound without external tracking using deep learning. *Med. Image Anal.* **48**, pp. 187–202 (2018)
17. Rahmatullah, B., Papageorghiou, A., Noble, J.A.: Automated Selection of Standardized Planes from Ultrasound Volume. In: *Machine Learning in Medical Imaging (MICCAI Workshop)*, pp. 35–42 (2011)
18. Salomon, L.J., Alfirevic, Z., Berghella, V., Bilardo, C., Hernandez-Andrade, E., Johnsen, S.L., Kalache, K., Leung, K.Y., Malinger, G., Munoz, H., Prefumo, F., Toi, A., Lee, W.: Practice guidelines for performance of the routine mid-trimester fetal ultrasound scan. *Ultrasound Obstet. Gynecol.* **37**(1), pp. 116–126 (2011)
19. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: MobileNetV2: Inverted Residuals and Linear Bottlenecks. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4510–4520 (2018)
20. Shah, S., Bellows, B.A., Adedipe, A.A., Totten, J.E., Backlund, B.H., Sajed, D.: Perceived barriers in the use of ultrasound in developing countries. *Crit. Ultrasound J.* **7**, 11 (2015)
21. Toporek, G., Wang, H., Balicki, M., Xie, H.: Autonomous image-based ultrasound probe positioning via deep learning. In: *Hamlyn Symposium on Medical Robotics* (2018)
22. Vilchis, A., Troccaz, J., Cinquin, P., Masuda, K., Pellissier, F.: A new robot architecture for tele-echography. *IEEE Trans. Robot. Autom.* **19**(5), pp. 922–926 (2003)
23. Wang, S., Housden, J., Noh, Y., Singh, D., Singh, A., Skelton, E., Matthew, J., Tan, C., Back, J., Lindenroth, L., Gomez, A., Toussaint, N., Zimmer, V., Knight, C., Fletcher, T., Lloyd, D., Simpson, J., Pasupathy, D., Liu, H., Althoefer, K., Hajnal, J., Razavi, R., Rhode, K.: Robotic-Assisted Ultrasound for Fetal Imaging: Evolution from Single-Arm to Dual-Arm System. In: *Towards Autonomous Robotic Systems (TAROS)*. pp. 27–38 (2019)
24. Yaqub, M., Kelly, B., Papageorghiou, A.T., Noble, J.A.: A Deep Learning Solution for Automatic Fetal Neurosonographic Diagnostic Plane Verification Using Clinical Standard Constraints. *Ultrasound Med. Biol.* **43**(12), pp. 2925–2933 (2017)