



City Research Online

City, University of London Institutional Repository

Citation: Sendra-Balcells, C., Campello, V. M., Torrents-Barrena, J., Ahmed, Y., Elattar, M., Ohene-Botwe, B., Nyangulu, P., Stones, W., Ammar, M., Benamer, L. N., et al (2023). Generalisability of fetal ultrasound deep learning models to low-resource imaging settings in five African countries. *Scientific Reports*, 13(1), 2728. doi: 10.1038/s41598-023-29490-3

This is the published version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/29758/>

Link to published version: <https://doi.org/10.1038/s41598-023-29490-3>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk



OPEN

Generalisability of fetal ultrasound deep learning models to low-resource imaging settings in five African countries

Carla Sendra-Balcells^{1✉}, Víctor M. Campello^{1✉}, Jordina Torrents-Barrena², Yahya Ali Ahmed³, Mustafa Elattar^{4,5}, Benard Ohene-Botwe^{6,7}, Pempho Nyangulu⁸, William Stones⁸, Mohammed Ammar⁹, Lamy Nawal Benamer¹⁰, Harriet Nalubega Kisémba¹¹, Senai Goitom Sereke¹², Sikolia Z. Wanyonyi¹³, Marleen Temmerman¹⁴, Eduard Gratacós^{15,16,17}, Elisenda Bonet^{15,18}, Elisenda Eixarch^{15,16,17}, Kamil Mikolaj¹⁹, Martin Grønnebak Tolsgaard¹⁹ & Karim Lekadir¹

Most artificial intelligence (AI) research and innovations have concentrated in high-income countries, where imaging data, IT infrastructures and clinical expertise are plentiful. However, slower progress has been made in limited-resource environments where medical imaging is needed. For example, in Sub-Saharan Africa, the rate of perinatal mortality is very high due to limited access to antenatal screening. In these countries, AI models could be implemented to help clinicians acquire fetal ultrasound planes for the diagnosis of fetal abnormalities. So far, deep learning models have been proposed to identify standard fetal planes, but there is no evidence of their ability to generalise in centres with low resources, i.e. with limited access to high-end ultrasound equipment and ultrasound data. This work investigates for the first time different strategies to reduce the domain-shift effect arising from a fetal plane classification model trained on one clinical centre with high-resource settings and transferred to a new centre with low-resource settings. To that end, a classifier trained with 1792 patients from Spain is first evaluated on a new centre in Denmark in optimal conditions with 1008 patients and is later optimised to reach the same performance in five African centres (Egypt, Algeria, Uganda, Ghana and Malawi) with 25 patients each. The results show that a transfer learning approach for domain adaptation can be a solution to integrate small-size African samples with existing large-scale databases in developed countries. In particular, the model can be re-aligned and optimised to boost the performance on African populations by increasing the recall to 0.92 ± 0.04 and at the same time maintaining a high precision across centres. This framework shows promise for building

¹Departament de Matemàtiques i Informàtica, Universitat de Barcelona, Barcelona, Spain. ²HP Inc., Barcelona, Spain. ³Obstetrics and Gynecology Department, School of Medicine, Suez University, Suez, Egypt. ⁴Medical Imaging and Image Processing, Center of Informatics Science, Nile University, Sheikh Zayed City, Egypt. ⁵Research and Development Division, Intixel, Cairo, Egypt. ⁶Department of Radiography, School of Biomedical and Allied Health Sciences, College of Health Sciences, University of Ghana, Accra, Ghana. ⁷Division of Midwifery and Radiography, School of Health and Psychological Sciences, University of London, London, UK. ⁸Kamuzu University of Health Sciences, Blantyre, Malawi. ⁹Department of Electrical Engineering Systems, Laboratory of Engineering System and Telecommunication, University of M'Hamed Bougara Boumerdes, Algiers, Algeria. ¹⁰Obstetrics and Gynecology Department, School of Medicine, Algiers University, Algiers, Algeria. ¹¹Department of Radiology, Mulago National Referral and Teaching Hospital, Kampala, Uganda. ¹²Department of Radiology and Radiotherapy, School of Medicine, Makerere University College of Health Sciences, Kampala, Uganda. ¹³Department of Obstetrics and Gynaecology, Aga Khan University Hospital, 3rd Parklands Avenue, Nairobi, Kenya. ¹⁴Centre of Excellence in Women and Child Health, Aga Khan University, Nairobi, Kenya. ¹⁵BCNatal Fetal Medicine Research Center, Hospital Clínic and Hospital Sant Joan de Déu, Universitat de Barcelona, Barcelona, Spain. ¹⁶Institut d'Investigacions Biomèdiques August Pi i Sunyer (IDIBAPS), Barcelona, Spain. ¹⁷Centre for Biomedical Research on Rare Diseases (CIBER-ER), Barcelona, Spain. ¹⁸Barcelona Tech, Universitat Politècnica de Catalunya, Barcelona, Spain. ¹⁹Copenhagen Academy for Medical Education and Simulation and Department of Obstetrics, Rigshospitalet, Copenhagen, Denmark. ✉email: carla.sendra@ub.edu; victor.campello@ub.edu

new AI models generalisable across clinical centres with limited data acquired in challenging and heterogeneous conditions and calls for further research to develop new solutions for the usability of AI in countries with fewer resources and, consequently, in higher need of clinical support.

Abbreviations

AI	Artificial Intelligence
CNN	Convolutional Neural Network
GE	General Electric

In the age of Big Data and digital health, great excitement has been generated around the extraordinary opportunities that artificial intelligence (AI) may offer in tomorrow's healthcare. In particular, deep learning-based methods have shown great promise for the analysis of complex biomedical data. In medical imaging, deep learning has already made an impact in a wide range of applications such as cardiology¹, cancer² or brain imaging³ among many others^{4,5}. However, existing developments have been mostly focused on applications in high-resource settings, where there is greater access to large clinical imaging datasets for training deep learning networks. Applications in environments with limited resources, such as countries across the African continent, have been scarce. Thus, there is a risk of increasing inequalities in global health due to the disparities in the application of AI in medical imaging.

Paradoxically, in places with important medical imaging needs the development of AI is at its lowest level. For example, in Sub-Saharan Africa, there is a very high level of neonatal mortality (i.e. 27 deaths per 1000 births in 2019, compared to an average rate of 3.4 per 1000 in the European Union⁶) and a very high burden of stillbirths⁷, but the access to antenatal screening is limited, especially in rural Africa. A statistical study from the World Health Organisation highlights a critical shortage of trained obstetricians across African countries⁸. However, the diagnosis of fetal abnormalities in clinical practice requires clinical expertise to acquire and interpret ultrasound images of the fetus.

At the same time, the literature is abundant with AI developments and deep learning implementations to facilitate the tasks of fetal ultrasound scanning, quantification and diagnosis^{9–14}. For example, several techniques have been developed to help clinicians acquire fetal ultrasound planes that are adequate for further quantification, i.e. that cover specific structures such as the fetal head, fetal femur, fetal abdomen and fetal thorax. Baumgartner et al.¹⁵ were the first to identify and localise 13 fetal standard planes in real time during 2D ultrasound mid-pregnancy examinations using convolutional neural networks (CNNs). Later, similarly but in a retrospective manner, Maraci et al.¹⁶ automated the task of detecting the fetal presentation and heartbeat from a pre-defined free-hand ultrasound video with support vector machines and a conditional random field model instead of CNNs, due to the limited number of samples available for their study. Alternative approaches involved obtaining standard planes from 3D ultrasound (US) volumes to learn the mapping between the 2D plane and the transformation required to move towards the standard plane within the 3D volume¹⁷. More recently, Burgos-Artizzu et al.¹⁸ evaluated the maturity of state-of-the-art CNNs to automatically classify 2D maternal fetal US and released a large open-source dataset to promote further research on the matter. While some AI tools are being implemented to recognise multiple views in fetal imaging¹⁹, others focus their work on specific anatomical structures^{20–22}.

Despite the big amount of research studies, the proposed techniques have been trained and/or tested with fetal ultrasound datasets from high-income countries, such as the UK^{15–17}, China²¹, United States¹⁹ or Spain^{18,20}. There is no evidence of the ability of these deep learning models to generalise in centres with low resources, i.e. with lower image quality and a limited number of data to calibrate the models. To our knowledge, only one study assessed the applicability of a fetal ultrasound plane classifier across countries, but exclusively in high-resource imaging settings: a classifier was trained in the UK and tested successfully in Denmark²³.

This paper presents the very first study investigating the transferability of AI models trained in high-income settings and their applicability to process images acquired in low-income settings. Specifically, we build fetal ultrasound classifiers using large-scale datasets from Spain and assess their performance when applied to fetal ultrasounds acquired in five different African countries with differences in scanners, populations and contexts. Furthermore, transfer learning, an effective strategy to adapt a pre-trained model to a new domain²⁴, is implemented and assessed under real-world conditions, i.e. based on very small size samples of new images from the local sites, to assess the resources and efforts needed to re-calibrate the original model to obtain high-performance in each local clinical centre. The main contributions of this work are summarised as follows:

- The collection and annotation of a new multi-centre fetal ultrasound dataset comprised of 120 patients from five African countries, namely Algeria, Egypt, Malawi, Uganda and Ghana. This dataset is the most diverse ultrasound dataset acquired to our knowledge and represents real clinical scenarios across low-resource imaging settings. As a result, the ultrasound appearance, both globally and locally, exhibits marked differences as clearly illustrated in Fig. 1.
- Application and evaluation of a fetal plane classification model trained in high-income settings within a diversity of clinical centres in low-resource imaging settings. To that end, a deep learning model is built using a large and diverse European dataset from Spain (12,400 images from two hospitals)¹⁸ and externally evaluated in two different settings: (1) with fetal ultrasounds from four centres located in another European country, namely Denmark²³, to assess generalisability under similar settings and resources; and (2) with the fetal ultrasound images acquired from five African countries to assess generalisability under different conditions and contexts.

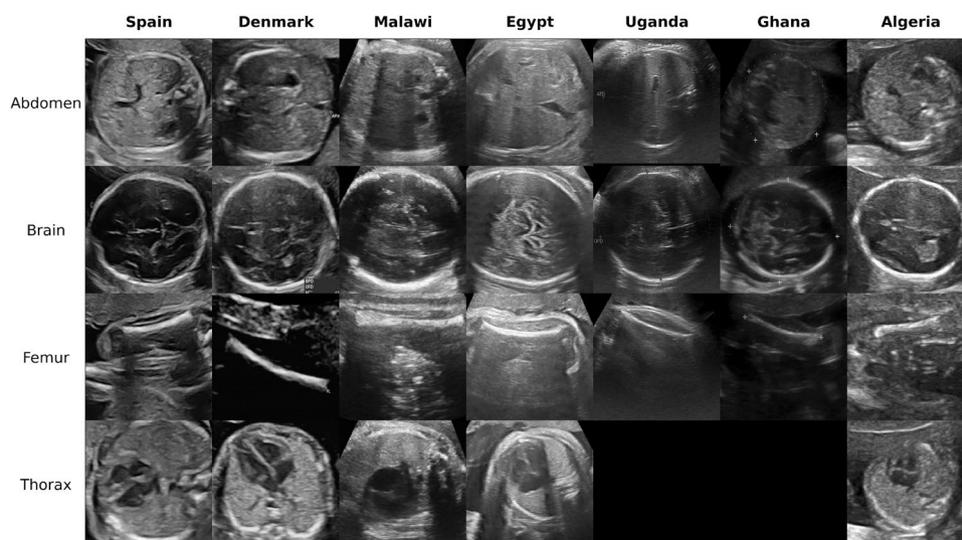


Figure 1. Image examples of the maternal-fetal US categories from our multi-centre dataset.

- Implementation of adaptive learning approaches to optimise the scalability of the classifiers in settings with distinct imaging conditions and lower image quality. This includes the estimation of the minimal number of samples required for calibrating the models to reach maximal performance for standard fetal plane classification in each of the local African centres. As a result, realistic strategies are derived for the scalability and applicability of existing deep learning models across new centres and imaging settings with minimal effort.

Methods

Datasets. The data used in this study and described in Table 1, consist of seven datasets acquired in distinct clinical centres. The first is a publicly available dataset from two centres in Spain¹⁸. The second has been acquired in four different centres in Denmark. Finally, the last five datasets contain US samples from five African countries (Malawi, Algeria, Uganda, Ghana and Egypt). All datasets were fully anonymised according to institutional guidelines. The Spanish dataset is publicly available, so no approval was necessary for this work. Subjects in Denmark, Ghana and Algeria signed informed consents and the studies were conducted under the approval of the ethics committee from each scan centre. The committee responsible in Denmark was the Danish Patient Safety Authority and the Danish Data Protection Agency. For Ghana, the Ethics and Protocol Review Committee of the School of Biomedical & Allied Sciences, University of Ghana. And for Algeria, the Ethics Committee of the Kouba Hospital, Algeria. For Malawi, Egypt and Uganda datasets, no individual consent was necessary according to national guidelines since no identifiable data was collected. The subjects have been scanned using a range of scanner US vendors by GE Medical Systems, Siemens, Edan Instruments, Shenzhen Mindray Bio-Medical Electronics and Aloka. Moreover, the frequency range of the curved transducer used during the acquisition ranges from 3 to 7.5 MHz and the screening was performed during the 2nd and/or 3rd trimester of the pregnancy period. Finally, clinical experts or research technicians who received extensive training in US classifica-

Country	Vendors	Type of transducer	Freq range (MHz)	Name of clinical centre	Trimester pregnancy
Spain	Voluson E6, Voluson S8 and Voluson S10 (GE Medical Systems, Zipf, Austria), and Aloka (Aloka CO., LTD.)	Curvilinear transducer	3 to 7.5	Hospital Clinic and Hospital Sant Joan de Déu	2nd and 3rd
Denmark	Voluson E6, Voluson S8 and Voluson S10 (GE Medical Systems, Zipf, Austria)	Curvilinear transducer	3 to 7.5	Copenhagen University Hospital Rigshospitalet, Hvidovre Hospital, Herlev Hospital and Nordsjællands Hospital Hillerød.	2nd and 3rd
Malawi	Mindray DC-N2 (Shenzhen Mindray Bio-Medical Electronics Co., Ltd, China/Germany)	Curvilinear transducer	3.5	Queen Elizabeth Central Hospital	2nd and 3rd
Egypt	Voluson P8 (GE Medical Systems, Zipf, Austria)	Curvilinear transducer	7	Sayedaty Center	2nd
Uganda	ACUSON X600 (Siemens)	Curvilinear transducer	3 to 7.5	Mulago National Referral Hospital	3rd
Ghana	EDAN DUS 60 (Edan Instruments, Inc., Shenzhen, China)	Curvilinear transducer	3.5 to 5	KBTH Polyclinic (Accra)	2nd and 3rd
Algeria	Voluson S8 (GE Medical Systems, Zipf, Austria)	Curvilinear transducer	3 to 7.5	EPH Kouba and Clinique Des Lilas	2nd and 3rd

Table 1. Details about the seven datasets considered in the study.

tion were asked to classify all images into femur, thorax, head and abdomen. The other less common planes were classified into the 'other' category. However, the five African datasets lack the 'other' category, and two of them do not contain samples of the thorax, representing a real clinical scenario in low-resource settings. Details about the number of samples used in each case are summarised in Table 2.

European datasets. The first large and open-source dataset with 1792 patients was collected at BCNatal¹⁸, an institution with two associated hospitals (Hospital Clínic and Hospital Sant Joan de Déu, Barcelona, Spain). All pregnant women attending routine pregnancy screening during the 2nd and 3rd trimesters between October 2018 and April 2019 were included in the study. Images were acquired using six different US machines by different operators with similar experience. The US machines used were three Voluson E6, one Voluson S8, one Voluson S10 (GE Medical Systems) and one Aloka. Images were taken using a curved transducer with a frequency range between 3 and 7.5 MHz. Operators were instructed to avoid using any type of post-processing or artefacts such as smoothing, noise, pointers or callipers when possible.

The second big dataset of this study was composed of 2nd and 3rd trimester US scans obtained from four fetal medicine centres in Denmark (Copenhagen University Hospital Rigshospitalet, Hvidovre Hospital, Herlev Hospital and Nordsjællands Hospital Hillerød). The scans were completed in the period between 2009 and 2017 and were obtained using three GE machines: Voluson E7, Voluson E8 and Voluson E10. Images were taken using a curved transducer with a frequency range between 3 and 7.5 MHz and no post-processing or artefacts were applied.

African datasets. Five African datasets were collected specifically for this study between November 2021 and February 2022. Each dataset included 25 patients and the images were not processed after the acquisition.

The first African dataset was acquired at Queen Elizabeth Central Hospital in Malawi using a Mindary DC-N2 US machine. The 2nd and 3rd trimester US samples were obtained between December and February using a curved transducer with a 3.5 MHz frequency.

The second small-size dataset was obtained at Sayedaty centre in Egypt using a Voluson P8 (GE), a GE Voluson series ultrasound with less image quality. The 2nd-trimester ultrasound scans were acquired between November and December using a curved transducer with a 7 MHz frequency.

The third African dataset consists of images from pregnant women during the 3rd trimester from the Mulago National Referral Hospital in Uganda. The US images were obtained in December 2021 with a Siemens US machine and using a curved transducer with a frequency range between 3 and 7.5 MHz.

The fourth dataset was acquired with an EDAN DUS 60 US machine and comprises samples from women in their 2nd or 3rd trimester pregnancy period. The images were acquired at the KBTH Polyclinic Centre, Accra, Ghana, in December. Operators used a curved transducer with a frequency range between 3.5 and 5 MHz.

The last dataset was acquired between November 2021 and December 2021 at the EPH Kouba and Clinique Des Lilas centres in Algeria and using the same acquisition as in the European datasets, but with a lower-quality GE machine.

Data preparation and pre-processing. All patient data was fully anonymised by removing the header in the original image and was later stored in Portable Network Graphic (PNG) format without compression to avoid quality degradation. Since US images do not contain colour information, images were stored as grayscale bit-maps. Images with visual artefacts were flagged by the clinician or technician and excluded from the analysis. We cropped the region containing most of the field of view but excluded the vendor logo and ultrasound control indicators. Finally, a min-max normalization was used after resizing the image to 224 × 224 to keep the same intensity range in images from the same dataset.

Model architecture and implementation details. Based on the model benchmark presented by Burgos et al.¹⁸ for fetal plane classification, the best performing method, a Densenet-169²⁵, was implemented in this study. The Spanish dataset described previously was used as the training set for the aforementioned benchmark, achieving a 0.94 average class precision. The Adam optimiser was used for training the model with a constant learning rate of 10^{-3} and first and second moments of 0.9 and 0.999, respectively. At each iteration, a batch of size

Country	Abdomen	Brain	Femur	Thorax	Other	TOTAL	Patients
Spain	711	1781	1040	1718	4213	9463	1792
Denmark	771	635	844	291	0	2541	1008
Malawi	25	25	25	25	0	100	25
Egypt	25	25	25	25	0	100	25
Uganda	25	25	25	0	0	75	25
Ghana	25	25	25	0	0	75	25
Algeria	25	25	25	25	0	100	25

Table 2. Number of subjects for each of the four datasets used during the training, validation and testing phases.

24 was used to calculate the loss function (cross-entropy loss) and optimise the network parameters. To improve the learning of the model, data augmentation was implemented during training. At each batch, images were randomly flipped, cropped between 0 and 20%, translated from 0 to 10 pixels and randomly rotated up to an angle of 15 degrees. The code was developed in PyTorch²⁶ and the hardware used during training and inference consisted of an 8 GB NVIDIA GeForce RTX 2080 Ti GPU. The different datasets were split into 50% for training and 50% for testing because of the small sample size of the African datasets. Different patients were included in each split to avoid data leakage. During training, 20% of the studies were kept for validation for those datasets with a large sample size case, while this amount was increased to 50% in the small sample size case. The same class distribution as in the Spanish dataset was enforced in the remaining datasets (40% 'other' and 15% for each of the other planes: femur, thorax, abdomen and head) and the missing categories of the African datasets ('other' or thorax) were filled with samples from the Spanish dataset.

Ablation study. Several strategies were considered to study the model transferability from high to low-resource imaging settings. We classified them into single-centre and multi-centre. The first approach uses a single centre in high-resource settings and the second approach includes also samples from the target centre in low-resource settings.

1. **Single-centre:** We train the neural network with the Spanish dataset using a different amount of samples ($n = 125, 250, 500, 1000, 2000, 4000$). Then, we evaluate the classification performance on the new datasets to test the generalisability of the classification model either on the Danish dataset, acquired in high-resource settings or on the African datasets acquired in low-resource settings.
2. **Multi-centre:**

Combination: We train the neural network directly using two samples with varying sizes from two different datasets, the Spanish dataset ($n = 0, 125, 250, 500, 1000, 2000, 4000$) and each of the African datasets separately ($n = 25$). With this, we investigate the effect of the sample size of the first centre on the performance of the model in the second centre.

Transfer learning: We fine-tune the last 4 layers of a pre-trained neural network with a few US images from the new clinical site in Africa. To that end, we pre-train the neural network with the Spanish dataset using again a different amount of samples ($n = 0, 125, 250, 500, 1000, 2000, 4000$), being $n = 0$ the case in which a model pre-trained on the public ImageNet database for object detection and image classification is used. Then, the model is fine-tuned with the available US samples from each of the African centres. Additionally, we estimate the minimum number of patients ($p = 2, 4, 6, 8, 10, 12$) needed from the African centre during the fine-tuning to obtain the desired level of performance, compared to the case in which no pre-trained model is used.

Performance evaluation. The average accuracy metric is used for monitoring the evolution of the model training. Thus, the model that reaches the maximum accuracy on the validation data and the minimum validation loss is saved as the final model and is later assessed against the testing set. For all experiments and results, confusion matrices are calculated and saved. The area under the curve (AUC) score, a metric that did not influence the training of the model, is used to assess all the experiments. Since the classification involved multi-label outputs, the AUC was computed as the average of the scores for one-versus-rest comparisons, i.e. AUC scores were computed taking each label as the positive class and all the rest as negatives.

Results

Generalisation of the fetal plane classifier to high or low-resource settings using a single-centre approach. In the baseline experiment, we evaluated the generalisation of the model trained with the Spanish dataset for different amounts of samples. Figure 2 shows that a good performance is achieved on the Danish dataset (average AUC greater than 95%) when a sufficiently large amount of samples is used ($n > 500$). Although the Spanish model has a good generalisation ability in another European dataset, the African clinical centres with low-resource settings manifest an inconsistent pattern and lower performance, especially when the number of training samples is scarce. In the case where the images are acquired using a similar protocol to that of the European datasets, such as Algeria (Table 1), the model shows a competitive performance (average AUC of $97.9 \pm 1.9\%$) when more than 1000 samples are used for training.

Transferability of the fetal plane classifier from high to low-resource settings using a multi-centre approach: combination or transfer learning. The first approach tested to improve the generalisability of the model from high to low-resource settings consisted of the combination of the Spanish dataset with the small-size datasets from each African centre. The AUC for these models separated by the African dataset used is shown in Fig. 3 (left plot) for a different number of training samples. The same procedure was followed for the second approach (right plot in Fig. 3), which is based on fine-tuning a model, pre-trained with data from Spain, using all patients in Malawi, Egypt, Uganda, Ghana and Algeria. When combining different datasets for training, the number of samples from Spain needed to reach the maximum performance in the African clinical centres is in most cases 500 or 1000. However, this number is not accurate and the good performance of the model is highly dependent on its choice, making it challenging to find a good balance between both datasets. On the other hand, in the transfer learning approach, the performance of the new clinical centres increases progres-

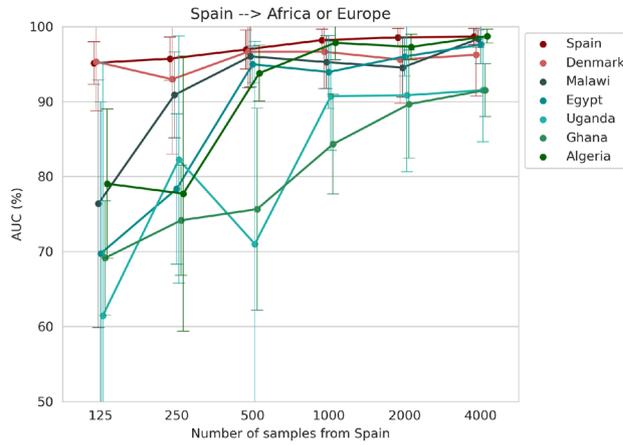


Figure 2. Generalisation performance of a model trained with different number of samples from the Spanish dataset evaluated in the same clinical centre, in a new centre in Europe (Denmark) and five datasets from Africa (Malawi, Egypt, Uganda, Ghana, Algeria).

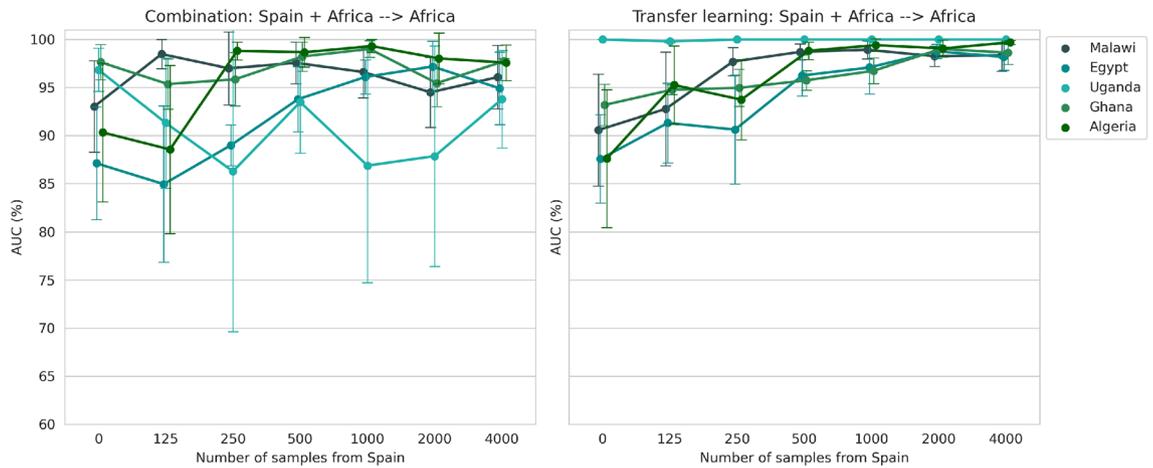


Figure 3. The AUC score as a result of combining the dataset acquired in Spain with the US samples acquired in each African centre independently or after fine-tuning the Spanish model, trained with varying sample size, with US images from each African centre independently.

sively with the number of cases from the Spanish dataset. The maximum AUC is obtained when using transfer learning (above 98% for all target datasets).

In terms of computational cost, Table 3 shows that training in the combination approach takes 8 times longer than when using transfer learning. Moreover, although the CPU and model storage is the same, the GPU resources needed are 3 times higher. In summary, the transfer learning approach is the most efficient considering the computational resources required and the final performance.

	Combination	Transfer learning
Inference time (s/epoch)	8	1
GPU memory (MB)	5683	1803
RAM memory (MB)	6000	6000
Model size (MB)	48.6	48.6
Batch size	24	24

Table 3. Summary of the computational resources needed for the combination ($n = 500$) and transfer learning approaches ($n = 4000$).

Given the potential of the transfer learning approach to result in a generalisable model we now study the optimal number of patients required to get a good performance. Using the model pre-trained with 4000 images from the Spanish dataset, we obtain in Fig. 4 the AUC for the different African datasets when using a different number of patients from the target domain. We compared this strategy with the case in which a big dataset is not available and so the model has to be trained using only the samples from the African centre. The results indicate that fine-tuning models trained on a large dataset with a small amount of target data is sufficient to reach a high classification performance. The generalisation performance increases progressively with the number of patients used from the target centre and reaches a maximum AUC when using around 4 or 6 patients. On the contrary, when fine-tuning is not used a lower performance is obtained. For example, in the case of Malawi, Egypt and Algeria, the maximum average AUC when training from scratch is 93.3%, 90.8% and 90.3%, respectively, while transfer learning increases the average performance to 99.2%, 98.8% and 99.8%. In Uganda and Ghana, the improvement given by transfer learning is smaller since the maximum baseline performance is already high: 99.5% and 97.6% versus 100% and 98.6% with fine-tuning, respectively. The higher baseline performance in Ghana and Uganda could be explained by the fact that in these two centres the classes used for testing are only 3 (abdomen, brain and femur). As observed in Fig. 5, the thorax plane is more frequently misclassified when fine-tuning is not used. In particular, most of the thorax examples are wrongly classified as brain or 'other'. Figure 5

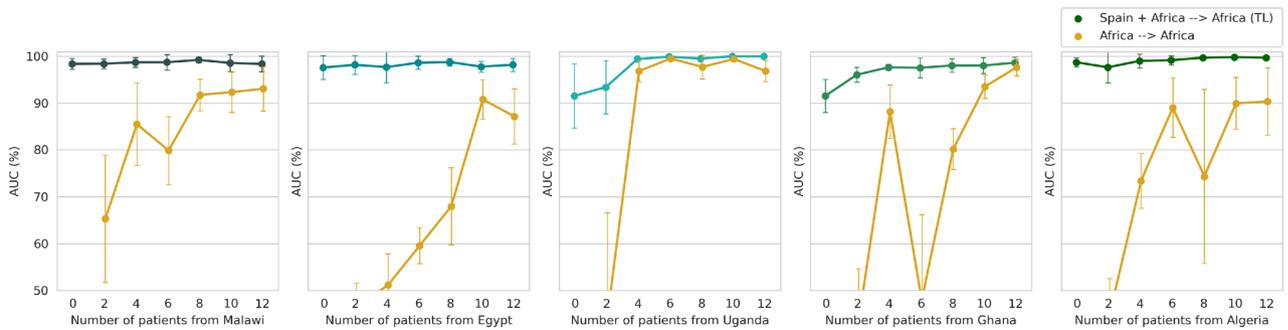


Figure 4. The evaluation of the model after being trained in Spain with all available samples and fine-tuned using a different number of patients from each African centre, as compared to the performance of the model directly trained with the African US images.

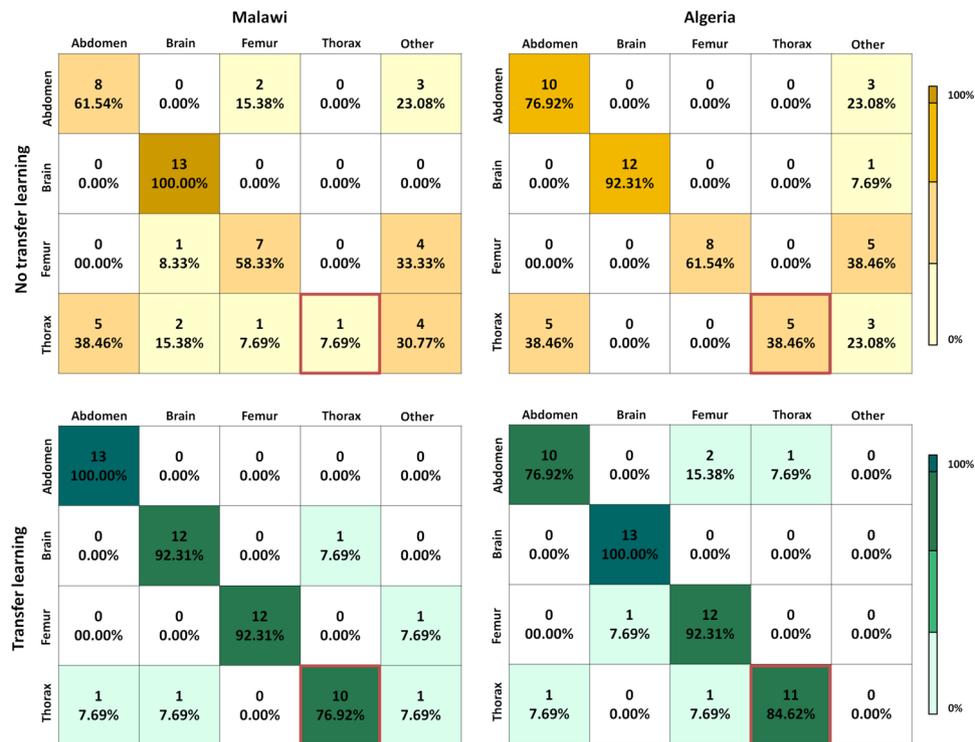


Figure 5. Results on common plane classification with and without transfer learning using 12 patients of each African centre in Malawi and Algeria.

also shows that when a model pre-trained in a large dataset is used and fine-tuned, the model can adapt to the new configuration of the new centre.

Finally, Fig. 6 quantitatively summarizes the capacity of transfer learning to adjust existing models to new centres with minimal effort and reduced sample size as compared to models trained from scratch with African datasets. Transfer learning achieves an improved average recall, by reducing the number of false negatives, while the average precision is stabilised and reaches a comparable value when compared to models trained with the Spanish dataset.

Discussion

An important goal for the adoption of deep learning models in the medical domain is to develop inclusive solutions that consider contexts with limited resources. However, to build robust predictive models, neural networks usually require a large number of parameters and large datasets, which are especially difficult to obtain in low-resource settings. For this reason, our work investigated the transferability of deep learning models, developed in centres with greater access to large clinical imaging datasets, to new centres with limited resources, where the generalisation ability is proved to be even more problematic.

Here, we used an existing large-scale data repository from Spain to build a model for US fetal plane classification that can be transferred to a new centre in the African continent. Our results showed that a model trained using a single-centre approach did not reach the desired performance when applied to a new centre in low-resource settings acquired in Malawi, Egypt, Algeria, Uganda or Ghana. In contrast, a multi-centre approach was able to bridge the gap between centres in Africa with less advanced equipment, and centres in Europe with high-quality images and similar acquisition. Specifically, we observed that by directly combining samples from both centres the performance was not as good and stable as with transfer learning, which was capable of effectively adapting the model to the new centre by only fine-tuning the last 4 layers of the model. Through shared representation learning, the neural network has favoured deep phenotypes that are common and consistent across populations/centres, while removing site-specific imaging patterns. With the optimisation of the final weights of the neural network, the objective is to maximise the predictive performance for the African ultrasound datasets, despite their small sample size.

Fine-tuning a model pre-trained with a large imaging dataset in Spain proved to be a potential solution to optimise the performance of a deep learning model in African settings by using small-size datasets and requiring minimal effort. This framework shows promise for generalisability across multi-centre African datasets with challenging and heterogeneous conditions. Moreover, it also encourages the community to conduct further research to develop highly generalisable solutions for the usability of AI in countries with fewer resources and consequently in higher need of clinical support.

One of the main limitations of the present study is that the datasets obtained in Africa are not complete, since they are less abundant and more difficult to compile. Future works should consider the evaluation of the model on the missing categories.

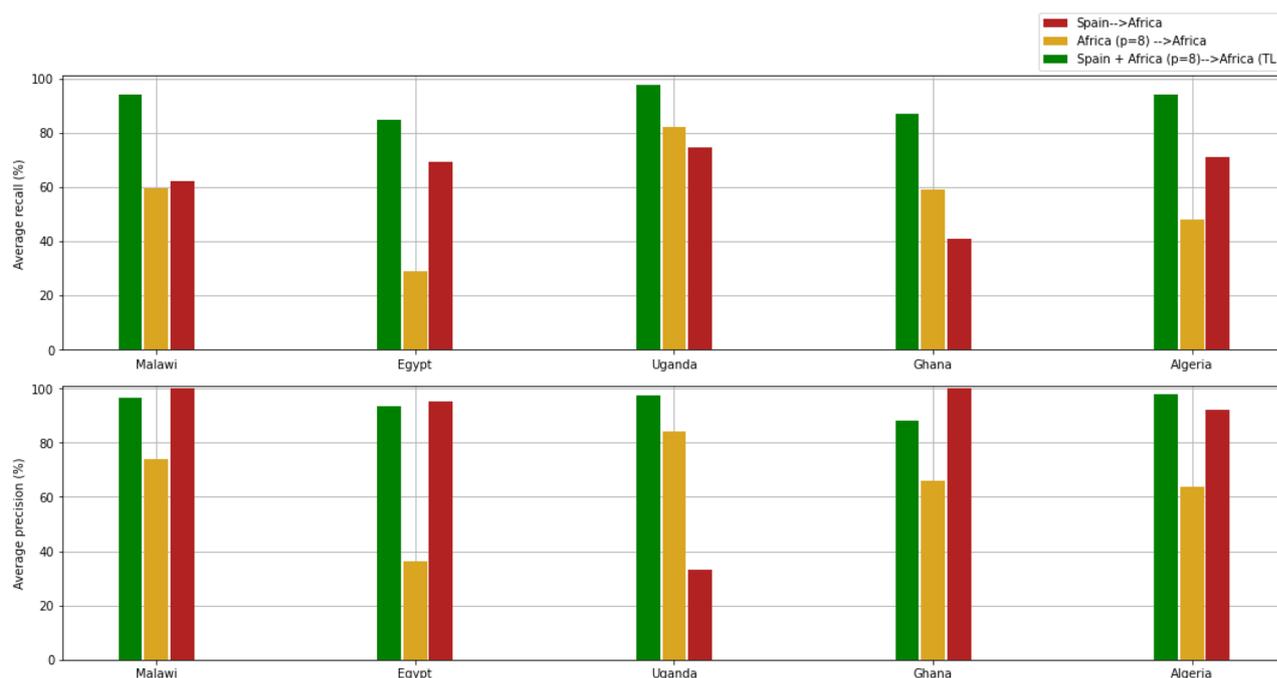


Figure 6. A summary of the performance achieved when only 8 patients from Egypt, Uganda, Ghana, Algeria and Malawi are used in terms of recall and accuracy. First, the model is directly trained with the African samples (yellow). Second, the model is trained using only the large Spanish dataset (red). Finally, the pre-trained Spanish model is fine-tuned with the new African samples (green).

Conclusions

This work was motivated by the need for building new AI models generalisable across clinical centres with limited data acquired in challenging and heterogeneous conditions. This framework has shown to be an opportunity for its widespread application to real maternal-fetal clinical settings, especially as a supporting tool in US plane recognition in low-income countries. We believe our work is an important first step in this direction and one that will encourage the development of more generalisable models based on transfer learning in centres with low-resource settings.

Data availability

The Spanish dataset is publicly available at zenodo.org/record/3904280. Access to the Danish dataset must be requested to the SONAI project members by contacting Martin G. Tolsgaard at martintolsgaard@gmail.com. The African datasets will be made available at Zenodo (<https://doi.org/10.5281/zenodo.7540448>) upon publication.

Received: 22 September 2022; Accepted: 6 February 2023

Published online: 15 February 2023

References

- Martin-Isla, C. *et al.* Image-based cardiac diagnosis with machine learning: A review. *Front. Cardiovasc. Med.* **7**, 1 (2020).
- Houssein, E. H., Emam, M. M., Ali, A. A. & Suganthan, P. N. Deep and machine learning techniques for medical imaging-based breast cancer: A comprehensive review. *Expert Syst. Appl.* **167**, 114161 (2021).
- Zhang, L., Wang, M., Liu, M. & Zhang, D. A survey on deep learning for neuroimaging-based brain disorder analysis. *Front. Neurosci.* **14**, 779 (2020).
- Zhou, S. K., Greenspan, H., Davatzikos, C., Duncan, J. S., Van Ginneken, B., Madabhushi, A., Prince, J. L., Rueckert, D., & Summers, R. M. A review of deep learning in medical imaging: Imaging traits, technology trends, case studies with progress highlights, and future promises. In *Proceedings of the IEEE* (2021).
- Aggarwal, R. *et al.* Diagnostic accuracy of deep learning in medical imaging: A systematic review and meta-analysis. *NPJ Digit. Med.* **4**(1), 1–23 (2021).
- WHO. Newborn mortality. <https://www.who.int/news-room/fact-sheets/detail/levels-and-trends-in-child-mortality-report-2021>
- De Bernis, L. *et al.* Stillbirths: Ending preventable deaths by 2030. *The Lancet* **387**(10019), 703–716 (2016).
- WHO. Number of licensed qualified obstetricians actively working. <https://www.who.int/data/gho/data/indicators/indicator-detail/GHO/number-of-licensed-qualified-obstetricians-actively-working>
- Torrents-Barrena, J. *et al.* Segmentation and classification in MRI and us fetal imaging: Recent trends and future prospects. *Med. Image Anal.* **51**, 61–88 (2019).
- Liu, S. *et al.* Deep learning in medical ultrasound analysis: A review. *Engineering* **5**(2), 261–275 (2019).
- Komatsu, M. *et al.* Detection of cardiac structural abnormalities in fetal ultrasound videos using deep learning. *Appl. Sci.* **11**(1), 371 (2021).
- Matthew, J. *et al.* Exploring a new paradigm for the fetal anomaly ultrasound scan: Artificial intelligence in real time. *Prenat. Diagn.* **42**(1), 49–59 (2022).
- Plotka, S. *et al.* Fetalnet: Multi-task deep learning framework for fetal ultrasound biometric measurements. In *International Conference on Neural Information Processing* 257–265 (Springer, 2021).
- Yin, J. *et al.* Ultrasonographic segmentation of fetal lung with deep learning. *J. Biosci. Med.* **9**(1), 146–153 (2021).
- Baumgartner, C. F. *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).
- 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**, 22–36 (2017).
- Li, Y., Khanal, B., Hou, B., Alansary, A., Cerrolaza, J. J., Sinclair, M., Matthew, J., Gupta, C., Knight, C., Kainz, B. & *et al.*, Standard plane detection in 3d fetal ultrasound using an iterative transformation network, in: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer, pp. 392–400 (2018)
- Burgos-Artizzu, X. P. *et al.* Evaluation of deep convolutional neural networks for automatic classification of common maternal fetal ultrasound planes. *Sci. Rep.* **10**(1), 1–12 (2020).
- Stephens, H. GE healthcare unveils AI-enabled ultrasound system. <https://axisimagingnews.com/radiology-products/imaging-equipment/ultrasound/ge-healthcare-ai-enabled-ultrasound-system>
- Montero, A., Bonet-Carne, E. & Burgos-Artizzu, X. P. Generative adversarial networks to improve fetal brain fine-grained plane classification. *Sensors* **21**(23), 7975 (2021).
- Wang, X. *et al.* Recognition of fetal facial ultrasound standard plane based on texture feature fusion. *Comput. Math. Methods Med.* <https://doi.org/10.1155/2021/6656942> (2021).
- Sushma, T. & Sriraam, N. Megha Arakeri, P. & Suresh, S. Classification of fetal heart ultrasound images for the detection of CHD. In *Innovative Data Communication Technologies and Application* 489–505 (Springer, 2021).
- Tolsgaard, M. *et al.* Does artificial intelligence for classifying ultrasound imaging generalize between different populations and contexts?. *Ultrasound Obstet. Gynecol.* <https://doi.org/10.1002/uog.23546> (2021).
- Kim, H. E. *et al.* Transfer learning for medical image classification: A literature review. *BMC Med. Imaging* **22**(1), 1–13 (2022).
- Huang, G., Liu, Z., Van Der Maaten, L., & Weinberger, K. Q. Densely connected convolutional networks, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 4700–4708 (2017).
- Paszke, A. *et al.* Pytorch: An imperative style, high-performance deep learning library. *Adv. Neural Inf. Process. Syst.* **32**, 1–12 (2019).

Author contributions

Design of the work: C.S.B., J.T.B., K.L. Image collection: Y.A.A., M.E., B.O.B., P.N., W.S., M.A., L.N.B., H.N.K., S.G.S., M.T. Interpretation of results: C.S.B., J.T.B., K.M., M.T., K.L. Manuscript draft: C.S.B., V.M.C., J.T.B., K.L. Manuscript review: all authors.

Funding

This work received funding from the European Union's 2020 research and innovation programme under Grant Agreement No. 825903 (euCanSHare project), as well as from the Spanish Ministry of Science, Innovation and

Universities under grant agreement RTI2018-099898-B-I00. Additionally, the research leading to these results has received funding from Cerebra Foundation for the Brain Injured Child (Carmarthen, Wales, UK).

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to C.S.-B. or V.M.C.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023