Automated detection and reacquisition of motion# degraded images in fetal HASTE imaging at 3 T

HHS Public Access

Author manuscript

Magn Reson Med. Author manuscript; available in PMC 2022 April 01.

Published in final edited form as:

Magn Reson Med. 2022 April ; 87(4): 1914–1922. doi:10.1002/mrm.29106.

Automated detection and reacquisition of motion-degraded images in fetal HASTE imaging at 3 T

Borjan Gagoski[1,2], **Junshen Xu**[3], **Paul Wighton**[4], **M. Dylan Tisdall**[5], **Robert Frost**[2,4], **Wei-Ching Lo**[6], **Polina Golland**[3,7], **Andre van der Kouwe**[2,4], **Elfar Adalsteinsson**[3,8], **P. Ellen Grant**[1,2]

[1]Fetal Neonatal Neuroimaging and Developmental Science Center, Boston Children's Hospital, Boston, Massachusetts, USA

[2]Department of Radiology, Harvard Medical School, Boston, Massachusetts, USA

[3]Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

[4]Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Charlestown, Massachusetts, USA

[5]Department of Radiology, Perelman School of Medicine, University of Pennsylvania, Philadelphia, Pennsylvania, USA

[6]Siemens Medical Solutions USA, Inc, Charlestown, Massachusetts, USA

[7]Computer Science and Artificial Intelligence Laboratory (CSAIL), Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

[8]Institute for Medical Engineering and Science, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

Abstract

Purpose: Fetal brain Magnetic Resonance Imaging suffers from unpredictable and unconstrained fetal motion that causes severe image artifacts even with half-Fourier single-shot fast spin echo (HASTE) readouts. This work presents the implementation of a closed-loop pipeline that automatically detects and reacquires HASTE images that were degraded by fetal motion without any human interaction.

Methods: A convolutional neural network that performs automatic image quality assessment (IQA) was run on an external GPU-equipped computer that was connected to the internal network of the MRI scanner. The modified HASTE pulse sequence sent each image to the external computer, where the IQA convolutional neural network evaluated it, and then the IQA score was sent back to the sequence. At the end of the HASTE stack, the IQA scores from all the slices were sorted, and only slices with the lowest scores (corresponding to the slices with worst image quality) were reacquired.

Correspondence: Borjan Gagoski, Fetal Neonatal Neuroimaging and Developmental Science Center, Department of Radiology, Harvard Medical School, Boston Children's Hospital, Boston, MA, USA. borjan.gagoski@childrens.harvard.edu.
Borjan Gagoski and Junshen Xu contributed equally to this work.

Results: The closed-loop HASTE acquisition framework was tested on 10 pregnant mothers, for a total of 73 acquisitions of our modified HASTE sequence. The IQA convolutional neural network, which was successfully employed by our modified sequence in real time, achieved an accuracy of 85.2% and area under the receiver operator characteristic of 0.899.

Conclusion: The proposed acquisition/reconstruction pipeline was shown to successfully identify and automatically reacquire *only* the motion degraded fetal brain HASTE slices in the prescribed stack. This minimizes the overall time spent on HASTE acquisitions by avoiding the need to repeat the entire stack if only few slices in the stack are motion-degraded.

Keywords

automated image quality assessment; fetal brain MRI; HASTE MRI with reacquisition

1 | INTRODUCTION

Fetal Magnetic Resonance Imaging (MRI) has become a well-established procedure that serves as an important adjunct to ultrasound when diagnostic doubts remain after the ultrasound exam.[1,2] The field of fetal brain MRI has been dominated by the use of single-shot encoding techniques such as half-Fourier acquisition single-shot turbo spin echo (HASTE) imaging, considered the workhorse of fetal MRI because of its ability to delineate cerebral anatomy and provide images with reduced artifacts due to fetal motion when compared to other T_2-weighted techniques.[3] Although typical readout lengths are ~500 ms (for voxel sizes = 1.3 × 1.3 × 3 mm^3, in-plane acceleration R = 2, partial Fourier = 5/8, echo spacing of ~6 ms), repetition times (TRs) can be more than 3 times longer (typical TR is ~1.8 s) due to specific absorption rate constraints. Therefore, it usually takes more than a minute to acquire a stack of ~40 slices to cover the entire fetal brain. Even with the relatively efficient readouts, unpredictable and random fetal motion still degrades image quality and introduces slice-to-slice variations of the imaging plane, resulting in double-oblique slices that are hard to interpret clinically. This significantly lengthens the fetal MRI sessions because the MRI technologist “chases” the fetal head in an attempt to obtain images in the standard 3 orthogonal planes, resulting in repetition of the entire HASTE stack if sufficient slices in the stack are motion-degraded or become double oblique. Thus, the current fetal imaging workflow depends on the experience of the MRI technologists performing the scan and their rapid assessment of the image stack quality in order to decide whether the entire stack should be repeated if enough slices are of bad quality.

In research settings, solutions such as the super-resolution slice-to-volume reconstruction[4–6] have been successfully used to generate high-quality fetal brain HASTE volumes, enabling further understanding of human brain development across gestation.[7–11] Slice-to-volume reconstruction methods treat the 3D volume reconstruction from 2D HASTE slices as an optimization problem and rely on oversampling of the fetal brain. Depending on the image quality, 6–12 orthogonal 2D HASTE stacks are typically required. Furthermore, most implementations require manual segmentation of the fetal head and visual selection of a good-quality HASTE stack as reference. Recent variants employ convolutional neural networks (CNN) to eliminate manual interaction by automating brain segmentation and outlier rejection[12] and even estimate the initial slice-to-volume transformations in real

time via CNN-based tracking.[13] Nevertheless, none of the methods to date address the unnecessary oversampling of HASTE stacks and image artifacts of the individual slices. Typically, if a certain number of slices in a given stack are of bad quality, the whole stack is repeated, resulting in unnecessary reacquisition of all good-quality slices and thus longer scan times. Lastly, and most importantly, the use of the slice-to-volume reconstruction methods still needs the deployment of computationally heavy postacquisition optimization, which precludes their routine use in clinical practice because very few sites have the infrastructure to reconstruct high-quality fetal brain HASTE volumes immediately after acquisition for diagnostic clinical read.

In this work, we designed and implemented a prototype closed-loop acquisition/reconstruction pipeline that automatically and without human interaction detects and reacquires *only* the HASTE slices that have been degraded by fetal motion, not the whole stack. This modified HASTE sequence makes use of a CNN developed recently by our team, which is trained to perform image quality assessment (IQA) on a fetal brain HASTE image in ~30 ms when run on a GPU-enabled computer. The feasibility of the proposed methods was tested on 10 pregnant mothers, showing that when there were image artifacts, the IQA CNN correctly identified the poor quality slices, and the sequence successfully reacquired only the slices with the worst IQA scores. Our modified sequence had 2 reacquisition modes. In the initially implemented mode, the number of slices to reacquire, N_{REACQ}, was defined by the user before the scan started; thus, a fixed number of slices were reacquired. In the second, more improved implementation, slices were reacquired only if they had an IQA score lower than a user-defined threshold, making N_{REACQ} dependent on how much the fetus moved during that particular scan. Both reacquisition modes improve the current standard clinical workflow by: (1) enabling savings in the overall scan times because only bad slices are reacquired, not the entire stack; and (2) providing objective means of determining which slices are nondiagnostic, as opposed to having the MR technologist make these decisions in real time.

2 | METHODS

2.1 | Convolutional neural network for image quality assessment

Recent literature[14,15] has demonstrated the potential of neural networks to perform image quality assessment on MR images. Inspired by these, we have developed a preliminary version of a CNN for fast fetal brain HASTE quality assessment[16,17] based on a pretrained ResNet architecture.[18] This initial work served as a starting point to further improve the robustness and reliability of the IQA CNN[19] applied in this work, and is discussed below.

We formulated the MR quality assessment as an image classification problem, with HASTE images being classified as “good” or “diagnostic,” “bad” or “nondiagnostic” (i.e., containing image artifacts or partial brain coverage), and those without the fetal brain within the acquired field-of-view (FOV). The CNN was initially trained using a dataset of ~11 thousand manually labeled HASTE images obtained from the clinical fetal MRI archives of Boston Children’s Hospital with appropriate institutional review board approval. The manual labeling was performed by a graduate student who was trained by an experienced neuroradiologist. Double oblique slices that showed no image artifacts were labeled as

"good." Figure 1A shows the distribution of our labeled dataset among the 3 categories, with 55%, 25%, and 20% of the images being classified as good, bad, and those without a brain within the FOV, respectively. Representative images from each category are also shown.

We trained our CNN using semi-supervised learning based on a mean teacher model[20] consisting of 2 networks: a student and a teacher network, both which use the same ResNet-34 architecture[18] as the backbone network. Using the manually labeled dataset, the CNN is trained to minimize the cross-entropy loss between the network prediction and the ground truth label, L_{cls}. Because manual labeling of MR images is difficult and time-consuming, a great feature of semi-supervised learning is its ability to make use of large unlabeled datasets to improve its accuracy. Therefore, to train our CNN, in addition to the ~11 thousand manually labeled HASTE images, we also included more than 200 thousand unlabeled fetal HASTE images obtained from the clinical fetal MRI archives at Boston Children's Hospital.

Consistency in our model is imposed by minimizing the Kullback-Leibler divergence between the outputs of the student and teacher networks, L_{con}. Because in fetal neuroimaging only the artifacts occurring within the fetal brain region of interest (ROI) affect diagnostic quality of the image, we introduce an ROI consistency loss to regularize the network so that it focuses on features within the brain ROI.[19] The ROI consistency loss, L_{roi}, is defined as the mean squared error between the features extracted from the brain ROI and the entire image. During training, the student network is updated by minimizing the total loss function $L = L_{cls} + L_{con} + L_{roi}$, whereas the teacher network is updated by recalculating its weights using moving average of the student network. Specifically, let w_t and w_s be the weights of the teacher and student networks, respectively. After each training step, w_t is updated to $\alpha \cdot w_t + (1 - \alpha) \cdot w_s$, where $\alpha = 0.994$ is the coefficient for the moving average. At the inference stage, the student network is used as the IQA network. A schematic of the proposed network model is given in Figure 1B.

All neural networks were implemented with PyTorch (version 1.4) and trained for 30 epochs using an Adam optimizer with a learning rate of 0.005. When evaluated on a separate testing dataset consisting of 1724 images, our semi-supervised CNN achieved 85.2% accuracy for ternary classification and area under the receiver operator characteristic for nondiagnostic images of 0.899, outperforming the model trained on supervised data only, as shown in Figure 1C. For a typical image matrix size of 256 × 256, the inference of the IQA network takes 30 ms on a GPU-enabled computer. We define $IQA_{score} = 1 - P_{bad}$, where P_{bad} is the probability of being a nondiagnostic image predicted by the network, resulting in a score close to 0 given to images with severe artifacts and a score of 1 given to those with excellent quality.

2.2 | Closed loop imaging framework for prospective IQA and reacquisition

The core of this work was the development and implementation of a modified HASTE sequence capable of performing *prospective* quality assessment of the fetal brain HASTE images and *automated* reacquisition of motion degraded slices. Schematic of the overall acquisition pipeline is shown in Figure 2. The IQA CNN described above was run on an external GPU-equipped (NVIDIA GTX 1050Ti, NVIDIA Corporate, Santa Clara, CA)

computer, which was connected to the internal network of the MRI scanner via 1 Gb Ethernet hub for efficient communication/feedback between this computer and the scanner's computer running our custom HASTE sequence and reconstruction. This setup is similar to those used for real-time neurofeedback in fMRI experiments[21–23] and prospective motion correction in neuroanatomical MRI.[24,25]

At the beginning of the TR, the HASTE readout is played, and immediately after that the reconstructed image is sent to the external computer. Python scripts receive this image, run the CNN, and an IQA score is computed in ~30 ms and sent to a buffer. At the scanner, a DICOM image is created and displayed, just like while running a standard HASTE scan. During the idle time due to specific absorption rate constraints, the sequence requests the IQA score, and Python scripts on the external computer send the score back to the sequence. This score, along with the slice index, is saved by the sequence for later use. At this point, one TR period is concluded. This workflow is repeated for all the slices in the prescribed stack. When the prescribed HASTE stack is completed, the sequence sorts all the IQA scores from all the slices and reacquires only the slices with the N_{REACQ} worst IQA scores. N_{REACQ} can be fixed or variable, depending on the chosen reacquisition mode. In the first mode, N_{REACQ} is defined by the user before the scan starts so a fixed number of slices are reacquired no matter how many of the slices in the stack were motion-degraded. In the second mode, only slices that exceed a user-defined IQA score are reacquired, which shortens the reacquisition loop in the case when there is little to no motion. In the case when the fetus moves a lot, this mode of operation enforces a user-defined upper limit on the number of slices to reacquire. For the purposes of this work, this upper limit was chosen to be one half of the number of prescribed slices.

Before replacement of slices, the IQA scores of the reacquired slices are compared to the IQA scores of the originally acquired slices to see which score is better. Then, on the scanner, 2 different image sets are generated as separate DICOM image series. The first one includes the original images from the prescribed stack (i.e., what the standard clinical HASTE scan outputs on the scanner), whereas in the second DICOM series HASTE slices that received low IQA scores are replaced with the N_{REACQ} reacquired slices. This enables direct and immediate comparison of the reacquired slices against the originally acquired ones right on the scanner’s console. This concludes a single run of the closed-loop HASTE acquisition framework.

2.3 | The MRI experiments

All MRI experiments were done on a 3 Tesla Siemens Skyra system (Siemens Healthcare, Erlangen, Germany) at Boston Children’s Hospital using the 30-channel body flex array in combination with the spine array (total of ~42–48 receive elements used). The closed-loop system was tested on 10 pregnant mothers who signed informed consent forms approved by Boston Children’s Hospital’s institutional review board. All fetuses scanned were normally developing, and their gestation age ranged between 23 and 38.7 weeks. The HASTE readouts had the following imaging parameters: TE = 119 ms, TR = 1.8 s, FOV = 33 × 33 cm^2, matrix size = 256 × 256, 1.3 × 1.3 × 3 mm^3 voxels, echo spacing = 5.81 ms, partial Fourier = 5/8, and in-plane GRAPPA acceleration of R_{GRAPPA} = 2. A total of 50

runs of our modified HASTE sequence were performed among all 10 fetuses using the first reacquisition mode for which the number of slices to reacquire was fixed. Further, a total of 23 runs were performed on 7 fetuses using the second mode of operation for which only slices below an IQA score of 0.9 were reacquired. The value of 0.9 for the threshold was chosen to balance the false positive and false negative rates for detecting low-quality images. Specifically, a threshold of 0.9 applied to our testing dataset yielded false positive and false negative rates of 0.15 and 0.14, respectively.

3 | RESULTS

Figure 3 demonstrates the feasibility of our methods *in utero* on a 29-gestational week fetus, showing the 2 DICOM image series provided as outputs on the scanner by our pipeline. The first one (Figure 3A) comprises the original HASTE stack, and the second one (Figure 3B) (created after the reacquisition loop has finished) substitutes the originally acquired, poor-quality images with the ones that were reacquired. The IQA scores are denoted below each image. With slices having an IQA scores of less than 0.9 flagged for reacquisition, the sequence correctly re-acquired the 7 slices satisfying this criterion (red-flagged in Figure 3A). In this case, the fetus happened to remain still during the reacquisition of the 7 slices, resulting in excellent image quality. With 25 slices needed to cover the whole brain and 7 reacquired slices, the overall scan time of this HASTE run was 57.6 s (TR = 1.8 s), only 12.6 s longer compared to the 45 s needed to acquire the whole stack without reacquisition.

Figure 4 shows more in utero results further demonstrating the feasibility of our methods. Two examples are shown from 2 separate whole-brain HASTE scans, prescribed axial (Figure 4A) and sagittal (Figure 4B) to the brain of a 29.1 gestational week fetus. In both cases, the number of slices to reacquire was fixed to $N_{REACQ} = 10$. Further, Figure 4C,D show examples of 2 HASTE scans on a 26.1 and 27.6 gestational week fetuses, respectively. For these 2 scans, our sequence reacquired only slices that had an IQA score lower than 0.9. All 4 examples show slices in the prescribed stacks that had compromised image quality, as reflected by the recorded low IQA scores. During the reacquisition process, image quality improved, as reflected by the high IQA scores of the reacquired slices.

4 | DISCUSSION

We have successfully developed and implemented a closed-loop acquisition/reconstruction pipeline that prospectively detects and automatically reacquires only poor-quality fetal brain HASTE images within the prescribed stack. Whereas prospective reacquisition of compromised MRI data has been established in postnatal imaging,[26–29] to our knowledge this is the first fetal MRI application of prospective reacquisition based on fast, automatic imaging quality assessment CNN. The methods demonstrated in this work improve the standard clinical fetal MRI workflow by: (1) enabling savings in the overall time the pregnant woman needs to be in the scanner because only bad slices are reacquired, not the entire stack; and (2) providing objective means of determining which slices are motion-degraded and therefore of nondiagnostic quality, eliminating the need for the MRI technologist to make these types of decisions in real time.

The current pipeline was tested on 10 fetuses in 2 reacquisition modes. The first one, in which the number of slices to be reacquired was fixed before the scan started, is less favored than the one in which only slices above a user-defined IQA score are reacquired. In this work, that threshold was chosen to be 0.9 as a balance between false positive and false negative rates for detecting low-quality images. Determining a potentially better criterion for choosing the optimal value for the IQA score threshold is out of the scope of this feasibility study.

Although providing improved workflow compared to the current way of conducting clinical fetal brain MRI, due to fetal motion the proposed methods do not necessarily reacquire the originally prescribed anatomical slices because the IQA CNN developed in this work does not have any information related to the position and orientation of the fetal head. Dealing with fetal head motion requires prospective correction of the slice positions based on the fetal head motion estimates. Prospective motion mitigation using low-resolution echo-planar volumetric navigators has been well developed and established technique *postnatally,* and has been shown to improve image quality in high-resolution volumetric structural acquisitions such as T_1-weighted MPRAGE and T_2-weighted SPACE.[26] Our future work will build upon our real-time feedback system, as well as our previous work embedding echo-planar volumetric navigators in the HASTE sequence,[30,31] to not only schedule reacquisition based on the IQA scores but also dynamically update the FOV coordinates before the acquisition of each HASTE slice in order to achieve complete, diagnostic-quality stacks without double obliquities. This also will open up the possibility for more complex out-of-plane acquisition schemes aimed at optimizing the slice-to-volume reconstruction methods.

5 | CONCLUSION

To our knowledge, this is the first time that prospective image quality assessment and automatic reacquisition of motion-degraded images has been employed in fetal brain HASTE imaging. Moreover, the implemented acquisition infrastructure serves as a basis for our future work and ultimate goal of developing an intelligent framework that “chases” the fetal head in real-time in order to obtain high-quality, diagnostic HASTE images of the entire fetal brain in the least amount of acquisition time.

ACKNOWLEDGMENT

The authors would like to thank Elizabeth Holland and Cindy Zhou for recruitng the pregnant mothers. This work would not be possible without the following funding sources: NIH R01EB017337, U01HD087211, R01HD100009, R00HD074649, R01HD099846, R01HD093578, R01HD085813, The NVIDIA Corporation, NIH NIBIB NAC P41EB015902.

Funding information

This study was supported by the National Institutes of Health (NIH) grants R01EB017337; U01HD087211; R01HD100009; R00HD074649; R01HD099846; R01HD093578; and R01HD085813. Also, by the National Institutes of Health (NIH)/the National Institute of Biomedical Imaging and Bioengineering (NIBIB)/The Neuroimaging Analysis Center (NAC) grant P41EB015902 and The NVIDIA Corporation

REFERENCES

1. Jarvis DA, Griffiths PD. Current state of MRI of the fetal brain in utero. J Magn Reson Imaging 2019;49:632–646. [PubMed: 30353990]
2. Griffiths PD, Bradburn M, Campbell MJ, et al. Use of MRI in the diagnosis of fetal brain abnormalities in utero (MERIDIAN): a multicentre, prospective cohort study. The Lancet 2017;389:538–546.
3. Gholipour A, Estroff JA, Barnewolt CE, et al. Fetal MRI: a technical update with educational aspirations. Concepts Magn Reson Part A Bridg Educ Res 2014;43:237–266. [PubMed: 26225129]
4. Kuklisova-Murgasova M, Quaghebeur G, Rutherford MA, Hajnal JV, Schnabel JA. Reconstruction of fetal brain MRI with intensity matching and complete outlier removal. Med Image Anal 2012;16:1550–1564. [PubMed: 22939612]
5. Kim K, Habas PA, Rousseau F, Glenn OA, Barkovich AJ, Studholme C. Intersection based motion correction of multislice MRI for 3-D in utero fetal brain image formation. IEEE Trans Med Imaging 2010;29:146–158. [PubMed: 19744911]
6. Pier DB, Gholipour A, Afacan O, et al. 3D super-resolution motion-corrected MRI: validation of fetal posterior fossa measurements. J Neuroimaging 2016;26:539–544. [PubMed: 26990618]
7. Gholipour A, Rollins CK, Velasco-Annis C, et al. A normative spatiotemporal MRI atlas of the fetal brain for automatic segmentation and analysis of early brain growth. Sci Rep 2017;7:476. [PubMed: 28352082]
8. Zhan J, Dinov ID, Li J, et al. Spatial-temporal atlas of human fetal brain development during the early second trimester. NeuroImage 2013;82:115–126. [PubMed: 23727529]
9. Chapman T, Matesan M, Weinberger E, Bulas DI. Digital atlas of fetal brain MRI. Pediatr Radiol 2010;40:153–162. [PubMed: 19774370]
10. Vasung L, Rollins CK, Yun HJ, et al. Quantitative in vivo MRI assessment of structural asymmetries and sexual dimorphism of transient fetal compartments in the human brain. Cereb Cortex 2020;30:1752–1767. [PubMed: 31602456]
11. Ortinau CM, Rollins CK, Gholipour A, et al. Early-emerging sulcal patterns are atypical in fetuses with congenital heart disease. Cereb Cortex 2019;29:3605–3616. [PubMed: 30272144]
12. Ebner M, Wang G, Li W, et al. An automated framework for localization, segmentation and super-resolution reconstruction of fetal brain MRI. NeuroImage 2020;206:116324. [PubMed: 31704293]
13. Singh A, Salehi SSM, Gholipour A. Deep predictive motion tracking in magnetic resonance imaging: application to fetal imaging. IEEE Trans Med Imaging 2020;39:3523–3534. [PubMed: 32746102]
14. Esses SJ, Lu X, Zhao T, et al. Automated image quality evaluation of T2-weighted liver MRI utilizing deep learning architecture. J Magn Reson Imaging 2018;47:723–728. [PubMed: 28577329]
15. Tajbakhsh N, Shin JY, Gurudu SR, et al. Convolutional neural networks for medical image analysis: full training or fine tuning? IEEE Trans Med Imaging 2016;35:1299–1312. [PubMed: 26978662]
16. Lala S Convolutional Neural Networks for Image Reconstruction and Image Quality Assessment of 2D Fetal Brain MRI. [Dissertation]. Massachusetts Institute of Technology; 2019. https://dspace.mit.edu/handle/1721.1/123171. Accessed June 7, 2019.
17. Lala S, Singh N, Gagoski B, et al. A deep learning approach for image quality assessment of fetal brain MRI. In Proceedings of the 27th Annual Meeting of ISMRM, Montréal, Québec, Canada, 2019, p. 839.
18. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2016.
19. Xu J, Lala S, Gagoski B, et al. Semi-supervised learning for fetal brain MRI quality assessment with ROI consistency. International Conference on Medical Image Computing and Computer Assisted Intervention – MICCAI 2020. Springer; 2020: 386–395.
20. Tarvainen A, Valpola H. Mean teachers are better role models: weight-averaged consistency targets improve semi-supervised deep learning results. Proceedings of the 31st International Conference on Neural Information Processing Systems, Curran Associates, Inc, 2017;30: 1195–1204.

21. Hinds O, Ghosh S, Thompson TW, et al. Computing momentto-moment BOLD activation for real-time neurofeedback. NeuroImage 2011;54:361–368. [PubMed: 20682350]
22. Lee J, Wighton P & Cauley SF et al. Application of simultaneous multi-slice (SMS) imaging to real-time fMRI for improved neurofeedback signal fidelity. Real-time Functional Imaging and Neurofeedback (rtFIN) Conference, Gainesville, FL, 2015. Abstract 29.
23. Wighton P, Karahanoglu FI, Tisdall MD, van der Kouwe AJW, Slice-by-slice prospective motion correction in EPI using a Markerless motion tracking system. In ISMRM Workshop on Motion Correction in MRI, Cape Town, South Africa, 2017.
24. Gilman J, Wighton P & Curran MT et al. Modulation of visual attention of blended faces and scenes in the FFA and PPA. Real-time Functional Imaging and Neurofeedback (rtFIN) Conference, Gainesville, FL, 2015. Abstract 54.
25. Frost R, Wighton P, Karahanoglu FI, et al. Markerless highfrequency prospective motion correction for neuroanatomical MRI. Magn Reson Med 2019;82:126–144. [PubMed: 30821010]
26. Tisdall MD, Hess AT, Reuter M, Meintjes EM, Fischl B, van der Kouwe AJW. Volumetric navigators for prospective motion correction and selective reacquisition in neuroanatomical MRI. Magn Reson Med 2012;68:389–399. [PubMed: 22213578]
27. White N, Roddey C, Shankaranarayanan A, et al. PROMO: real-time prospective motion correction in MRI using image-based tracking. Magn Reson Med 2010;63:91–105. [PubMed: 20027635]
28. Frost R, Biasiolli L, Li L, et al. Navigator-based reacquisition and estimation of motion-corrupted data: application to multi-echo spin echo for carotid wall MRI. Magn Reson Med 2020;83:2026–2041. [PubMed: 31697862]
29. Bogner W, Gagoski B, Hess AT, et al. 3D GABA imaging with real-time motion correction, shim update and reacquisition of adiabatic spiral MRSI. NeuroImage 2014;103:290–302. [PubMed: 25255945]
30. Gagoski B, McDaniel P & van der Kouwe AJW et al. HASTE imaging with EPI volumetric navigators for real-time fetal head motion detection. In Proceedings of the 24th Annual Meeting of ISMRM, Singapore, 2016. p. 4413.
31. McDaniel P, Gagoski B, Tisdall MD, et al. Quantification of fetal motion tracked with volumetric navigator MRI acquisitions. In Proceedings of the 23rd Annual Meeting of ISMRM, Toronto, Ontario, Canada, 2015. p. 2576.

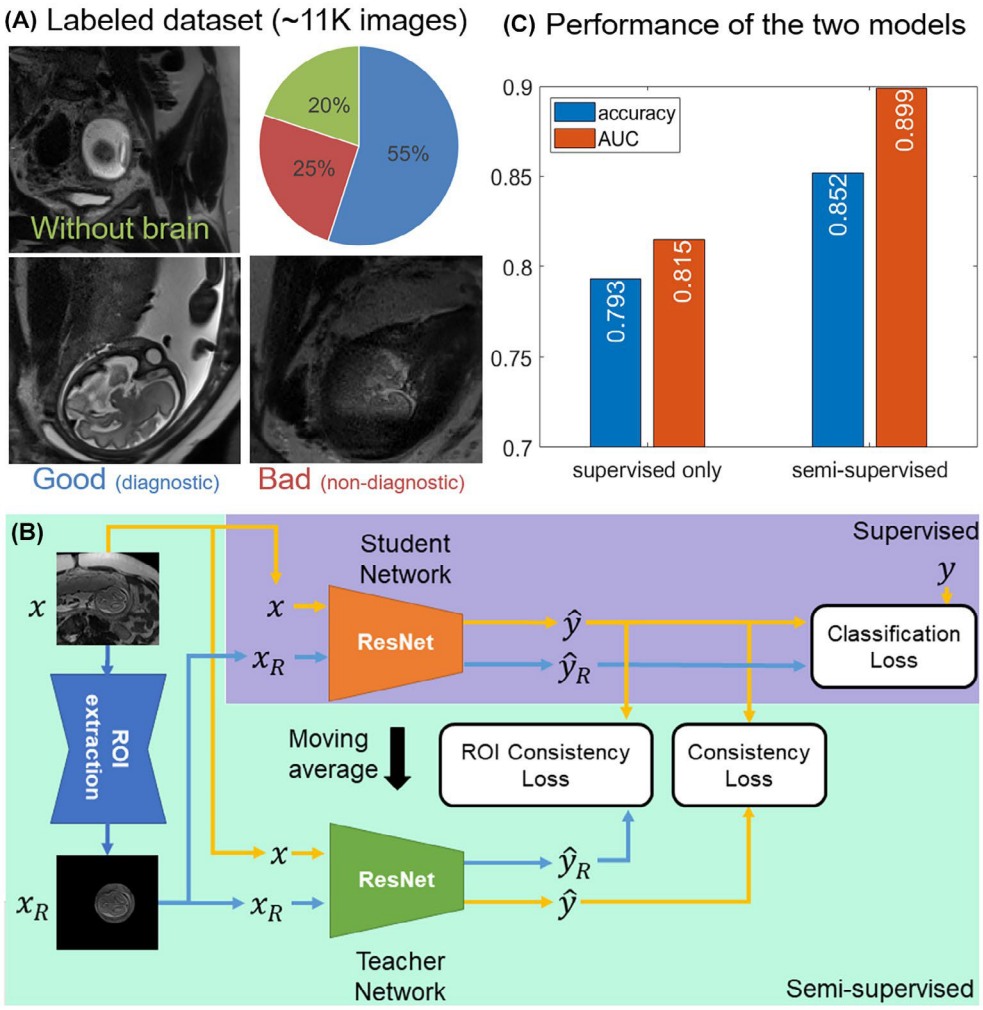

FIGURE 1.

(A) Distribution of the ~11 thousand HASTE-labeled fetal images into the 3 categories ("bad," "good," and "without a brain"), along with representative images for each category. An image without image artifacts was labeled as "good" even if it is a doubleoblique, as seen in the representative example. (B) Schematic of our semi-supervised learning methods for fetal brain HASTE image quality assessment. The model consists of a student and a teacher network, both of which are ResNet-34. The student network is trained to minimize the sum of classification loss, consistency loss, and ROI consistency loss, and the teacher network is updated as the moving average of the student network. After training, the student network is the one used as the IQA network, which generates the IQA scores. (C) Comparison between the performances of the model trained only using supervised data and our proposed semi-supervised learning method, showing the superiority of the latter. HASTE, half-Fourier acquisition single-shot turbo spin echo; IQA, image quality assessment; ROI, region of interest

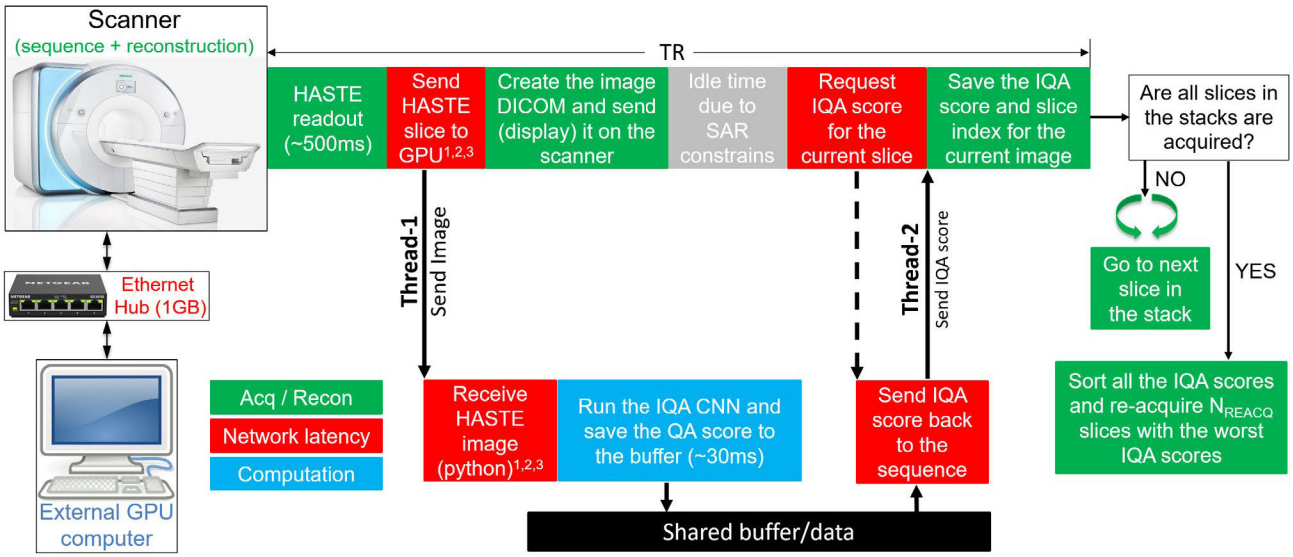

FIGURE 2.

Schematic of the closed-loop acquisition/reconstruction framework of the modified HASTE sequence capable of automatically detecting and reacquiring motion-degraded fetal brain HASTE images. The scanner's computer is connected with the external GPU computer via a 1GB ethernet hub. The GPU computer receives HASTE images in real time, calculates the IQA score for each of them, and sends this score back to the sequence upon request during runtime. Once all the slices of the prescribed stack are acquired, the sequence sorts all the IQA scores and automatically reacquires N_{REACQ} slices with the lowest scores. N_{REACQ}, number of slices to reacquire

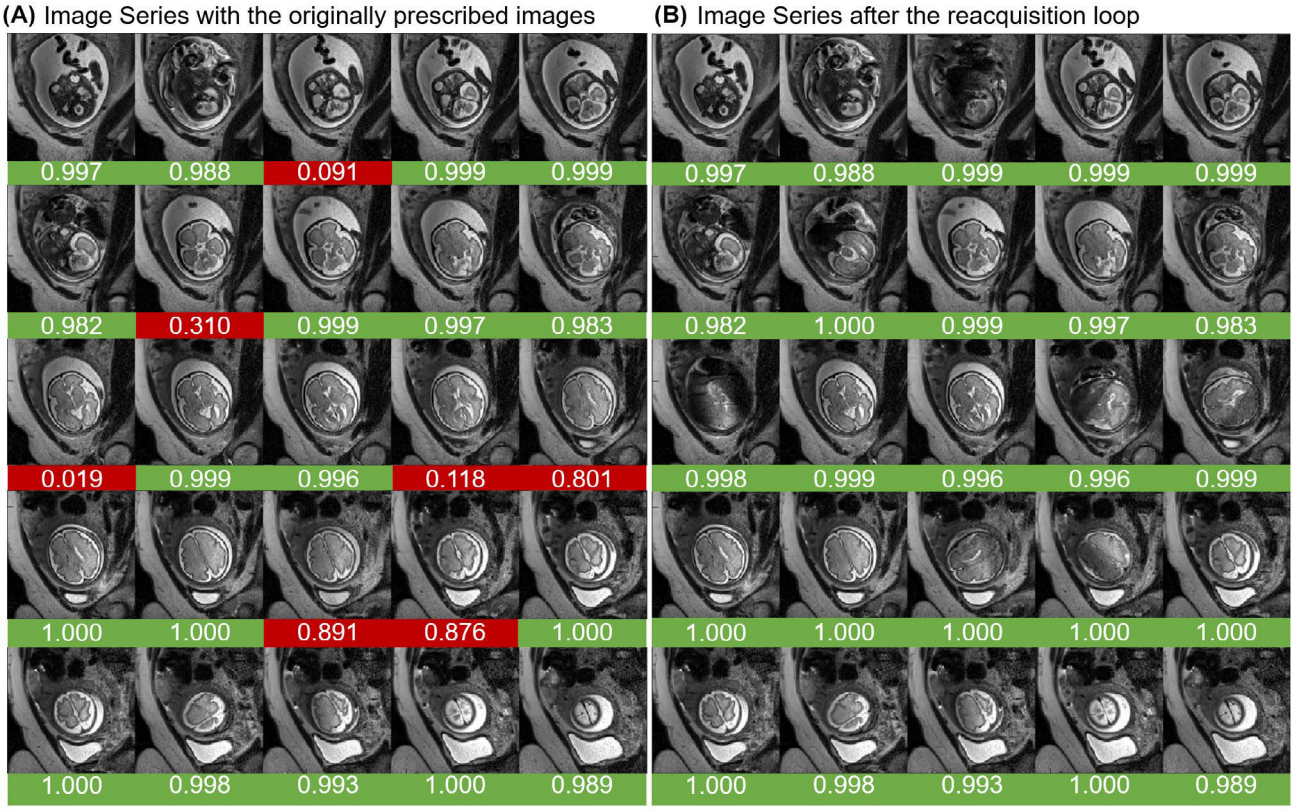

FIGURE 3.

In utero images from a 29 gestational week fetus using our modified HASTE sequence with the reacquisition mode set to reacquire only the slices that have an IQA score lower than 0.9. The IQA scores are denoted below each image. Our reconstruction pipeline provides 2 DICOM image series on the MRI console. The images shown in (A) are from the first image series, which correspond to 1 complete run over the prescribed stack, where we can see that 7 slices have met the reacquisition criterion (these are red-flagged). Images shown in (B) are from the second image series, which keeps all the slices that had an IQA score greater than 0.9 but now substitutes the original, degraded slices in (A) with the reacquired, improved images. We can see that, in this case, the 7 newly acquired slices were not degraded at the time of reacquisition and received high IQA scores. With 25 slices needed to cover the whole brain and 7 reacquired slices, the overall sequence time in (B) was 32 TRs, or 57.6 s compared to 45 s in (A) using a TR of 1.8 s

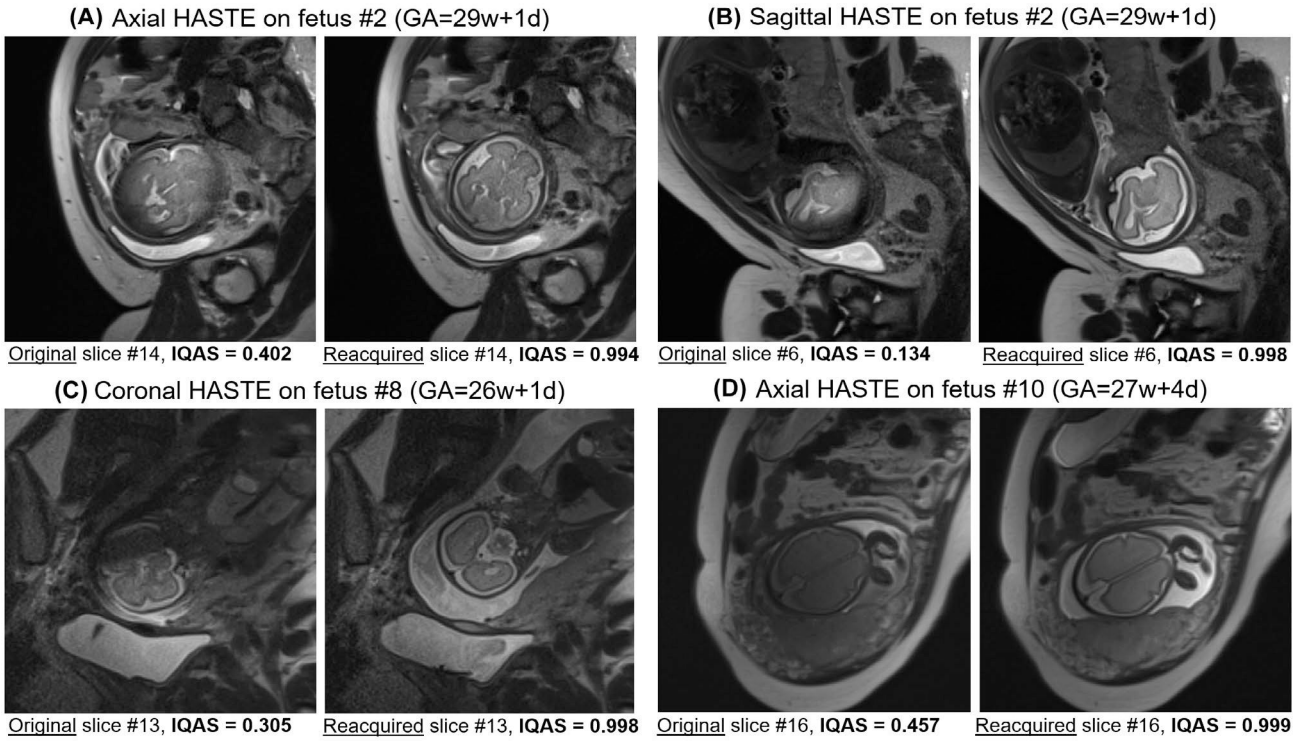

FIGURE 4.

Examples from 4 separate scans using our closed-loop system performed on 3 different fetuses. All scans were prescribed to cover the entire fetal brain. For the acquisitions in (A) and (B), the number of slices to reacquire was fixed before the start of the scan to $N_{REACQ} = 10$. For (C) and (D), only slices that had an IQA score lower than 0.9 were reacquired. The representative images shown demonstrate cases in which the original images are motion-degraded (and thus received a low IQA score), whereas the re-acquired images at the same slice locations suffered no motion degradation (and thus received a high IQA score)