

Comprehensive Evaluation of Unsupervised Image Enhancement for Volumetric Fetal Brain MRI

Yingqi Hao^{*1}Mingxuan Liu^{*1}Hongjia Yang¹Haoxiang Li¹Xuguang Bai¹Yi Liao²Haibo Qu²Qiyuan Tian^{†1}

HAO-YQ22@MAILS.TSINGHUA.EDU.CN

ARKTISX@FOXMAIL.COM

YANGHJ23@MAILS.TSINGHUA.EDU.CN

LIHAOXIA24@MAILS.TSINGHUA.EDU.CN

BXG21@MAILS.TSINGHUA.EDU.CN

CONNIE0064@126.COM

WINDOWSQHB@126.COM

QIYUANTIAN@TSINGHUA.EDU.CN

¹ School of Biomedical Engineering, Tsinghua University² Department of Radiology, West China Second University Hospital, Sichuan University**Editors:** Under Review for MIDL 2025

Abstract

MRI provides superior soft tissue contrast over ultrasound, making it essential for evaluating fetal brain development and pathology. However, clinical use of 2D thick-slice T2-weighted imaging remains constrained by motion-induced artifacts that degrade both image quality and subsequent quantitative analyses. While existing volumetric reconstruction pipelines (e.g., NiftyMIC, NeSVoR) incorporate motion correction, their outputs often retain residual noise due to a lack of post-reconstruction enhancement solutions. Although the recently proposed foundation model BME-X represents the first dedicated approach for fetal MRI enhancement, its generalizability to heterogeneous clinical datasets remains unproven. To bridge this gap, we conduct the first comprehensive comparison of BME-X and other unsupervised image enhancement methods on normal and pathological fetal brain MRI, based on tissue segmentation accuracy, tissue contrast t-score (TCT), lesion fidelity, and reader assessment. Results show that a pre-trained 3D convolutional variational autoencoder (VAE) achieves more effective enhancement compared to BME-X. Code and pre-trained weights are available at <https://github.com/yingqihao2022/FetalBrainEnhancement>.

Keywords: Fetal brain MRI, Image enhancement, BME-X, Foundation model, 3D VAE.

1. Introduction

Structural MRI’s non-invasive safety, superior soft tissue contrast, and radiation-free nature make it ideal for fetal brain characterization (Sun et al., 2024; Hao et al., 2025; Sanchez et al., 2023; Gholipour et al., 2011). However, 2D thick-slice T2-weighted acquisitions are vulnerable to motion interference from fetal movements and maternal respiration, causing blurring/ghosting that complicates neuroimaging analysis (Liu et al., 2025; Xu et al., 2020). While reconstruction tools like NiftyMIC (Ebner et al., 2019) and NeSVoR (Xu et al., 2023) can convert motion-corrupted 2D stacks into 3D volumes, persistent artifacts and noise degrade output quality (Sun et al., 2024).

* Contributed equally

† Corresponding author




Figure 1: Examples on a normal fetus (a) and a GMH-IVH fetus (b). The first row: axial slices of input/output images. The second row: residual maps (input - output). The third row: tissue segmentation results.

Enhancing volumetric fetal brain MRI faces two key challenges: the lack of noise-free ground truth, which limits supervised approaches, and the presence of non-traditional noise types, making unsupervised methods like Noise2Void (Krull et al., 2018), Zero-shot Noise2Noise (Mansour and Heckel, 2023), and Self2Self (Quan et al., 2020) less effective. Recently, Sun et al. introduced BME-X (Sun et al., 2024), the first successful model designed for this task, aiming to improve image quality by sharpening details and unifying tissue intensity ranges. While advertised as applicable to brain MRIs across all ages, their experiments focus on neonatal and adult cases, with limited validation on fetal images. To address the above problems, this study evaluates unsupervised image enhancement models for volumetric fetal brain MRI for the first time, including both normal and pathological cases. Evaluation includes tissue segmentation accuracy, the tissue contrast t-score (TCT) (Pizarro et al., 2016), lesion fidelity, and reader assessment. Results show that the pre-trained 3D convolutional VAE achieves state-of-the-art (SOTA) image enhancement, excelling in preserving brain structures and lesion fidelity.

2. Methods

2.1. Data Acquisition and Preprocessing

The study used MRI data from 442 pregnant women, with 432 normal fetal brains and 10 brains with Germinal Matrix Hemorrhage and Intraventricular Hemorrhage (GMH-IVH) lesions. 2D T2-weighted TSE imaging data were acquired in axial, coronal, and sagittal directions. The NeSVoR (Xu et al., 2023) method was used for slice-to-volume motion correction and 3D volumetric reconstruction.

2.2. Competing Methods

Five unsupervised image enhancement methods were implemented for fetal brain MRI: (1) 3D Convolutional VAE: Previous studies (Prakash et al., 2021) have explored direct denoising with VAE in microscopic imaging. The VAE model is designed to sample diverse solutions from trained posteriors by directly predicting per-pixel mean and variance, enabling the potential to learn right values. In this study, the MONAI framework (Pinaya et al., 2023) is utilized to implement a 3D VAE. (2) DAE: The denoising autoencoders (DAEs)




Figure 2: Reader assessment and TCT test results from different methods.




Figure 3: Enhanced results of the VAE model for brain volumes at various GAs (22w-38w).

eliminate noise through trained reconstruction of clean data from corrupted inputs. Our implementation adopts the DAE framework established in prior research (Hao et al., 2025). (3) BME-X: A recently proposed foundation model for brain MRI image enhancement (Sun et al., 2024). (4) BM4D: BM4D is an algorithm for attenuation of additive spatially correlated stationary (aka colored) Gaussian noise for volumetric data (Maggioni et al., 2013). (5) Self2Self: A self-supervised denoising technique using internal image statistics without external training data (Quan et al., 2020).

The VAE and DAE models were trained on 390 normal volumetric fetal brain MRI images. BME-X used its original pre-trained weights. Self2Self and BM4D are unsupervised single-image methods requiring no additional training data. The remaining 42 normal and 10 GMH-IVH fetal brain MRI volumes were reserved for testing.

3. Results and Conclusion

Figure 1 shows the enhancement results of normal and GMH-IVH fetal brain MRI images using different methods. Key observations include: (1) Residual maps reveal that only VAE effectively removed noise while preserving normal tissue contrast, while other methods changed tissue contrast during denoising, as confirmed by segmentation results (using a pretrained model (Fidon et al., 2024)); (2) Only VAE, BM4D, and Self2Self preserved hemorrhagic lesions in GMH-IVH brains, ensuring lesion fidelity. Figure 2 (a) presents reader assessment by an experienced radiologist (Y.L.), showing VAE’s balanced performance in signal-to-noise ratio (SNR) and tissue contrast. Figure 2 (b) illustrates TCT metrics (calculated following (Sun et al., 2024)), with VAE and DAE achieving top scores. Figure 3 illustrates VAE’s enhanced results across various gestational ages (GAs), demonstrating its generality. In conclusion, VAE is the SOTA volumetric fetal brain MRI enhancement model. Future work should explore VAE-based models to improve fetal brain image enhancement.

Acknowledgments

Funding was provided by the Tsinghua University Startup Fund, Dushi Program (grant number 20241080026) and Initiative Scientific Research Program.

References

- Michael Ebner, Guotai Wang, Wenqi Li, Michael Aertsen, Premal Amrishkumar Patel, Rosalind Aughwane, Andrew Melbourne, Tom Doel, Steven Dymarkowski, Paolo De Coppi, Anna Louise David, Jan A. Deprest, Sébastien Ourselin, and Tom Vercauteren. An automated framework for localization, segmentation and super-resolution reconstruction of fetal brain mri. *Neuroimage*, 206, 2019. URL <https://api.semanticscholar.org/CorpusID:207844623>.
- Lucas Fidon, Michael Aertsen, Florian Kofler, Andrea Bink, Anna L. David, Thomas Deprest, Doaa Emam, Frédéric Guffens, András Jakab, Gregor Kasprian, Patric Kienast, Andrew Melbourne, Bjoern Menze, Nada Mufti, Ivana Pogledic, Daniela Prayer, Marlene Stuempflen, Esther Van Elslander, Sébastien Ourselin, Jan Deprest, and Tom Vercauteren. A dempster-shafer approach to trustworthy ai with application to fetal brain mri segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(5):3784–3795, 2024. doi: 10.1109/TPAMI.2023.3346330.
- Ali Gholipour, Judy A. Estroff, Carol E. Barnewolt, Susan A. Connolly, and S. Warfield. Fetal brain volumetry through mri volumetric reconstruction and segmentation. *International Journal of Computer Assisted Radiology and Surgery*, 6:329–339, 2011. URL <https://api.semanticscholar.org/CorpusID:13042271>.
- Yingqi Hao, Mingxuan Liu, Juncheng Zhu, Hongjia Yang, Yi Liao, Haibo Qu, and Qiyuan Tian. Unsupervised Anomaly Detection for Fetal Brain MRI using Two-Stage Denoising Autoencoder (ω -DAE). In *2025 OHBM Annual Meeting*, Brisbane, Australia, June 2025. URL <https://hal.science/hal-04974207>.
- Alexander Krull, Tim-Oliver Buchholz, and Florian Jug. Noise2void - learning denoising from single noisy images. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2124–2132, 2018. URL <https://api.semanticscholar.org/CorpusID:53751136>.
- Mingxuan Liu, Haoxiang Li, Zihan Li, Hongjia Yang, Jialan Zheng, Haibo Qu, and Qiyuan Tian. Unsupervised Fetal Brain MRI Quality Assessment based on Orientation Prediction Uncertainty. In *2025 OHBM Annual Meeting*, Brisbane, Australia, June 2025. URL <https://hal.science/hal-04974115>.
- Matteo Maggioni, Vladimir Katkovnik, Karen Egiazarian, and Alessandro Foi. Nonlocal transform-domain filter for volumetric data denoising and reconstruction. *IEEE Transactions on Image Processing*, 22(1):119–133, 2013. doi: 10.1109/TIP.2012.2210725.
- Youssef Mansour and Reinhard Heckel. Zero-shot noise2noise: Efficient image denoising without any data. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recog-*

- nition (CVPR), pages 14018–14027, 2023. URL <https://api.semanticscholar.org/CorpusID:257631645>.
- Walter Hugo Lopez Pinaya, Mark S. Graham, Eric Kerfoot, Petru-Daniel Tudosiu, Jessica Dafflon, Virginia Fernandez, Pedro Sanchez, Julia Wolleb, Pedro F. Da Costa, Ashay Patel, Hyungjin Chung, Can Zhao, Wei Peng, Zelong Liu, Xueyan Mei, Oeslle Lucena, Jong-Chul Ye, Sotirios A. Tsaftaris, Prerna Dogra, Andrew Feng, Marc Modat, Parashkev Nachev, Sébastien Ourselin, and Manuel Jorge Cardoso. Generative ai for medical imaging: extending the monai framework. *ArXiv*, abs/2307.15208, 2023. URL <https://api.semanticscholar.org/CorpusID:260316368>.
- Ricardo Pizarro, Xi Cheng, Alan S. Barnett, Hervé Lemaître, Beth A. Verchinski, Aaron L. Goldman, Ena Xiao, Qian Luo, Karen Faith Berman, Joseph H. Callicott, Daniel R. Weinberger, and Venkata S. Mattay. Automated quality assessment of structural magnetic resonance brain images based on a supervised machine learning algorithm. *Frontiers in Neuroinformatics*, 10, 2016. URL <https://api.semanticscholar.org/CorpusID:18197826>.
- Mangal Prakash, Alexander Krull, and Florian Jug. Fully unsupervised diversity denoising with convolutional variational autoencoders. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=agHLCOBM5jP>.
- Yuhui Quan, Mingqin Chen, T. Pang, and Hui Ji. Self2self with dropout: Learning self-supervised denoising from single image. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1887–1895, 2020. URL <https://api.semanticscholar.org/CorpusID:219619080>.
- Thomas Sanchez, Oscar Esteban, Yvan Gomez, Elisenda Eixarch, and Meritxell Bach Cuadra. Fetmrqc: Automated quality control for fetal brain mri. In *PIPPI@MICCAI*, 2023. URL <https://api.semanticscholar.org/CorpusID:258078843>.
- Yue Sun, Limei Wang, Gang Li, Weili Lin, and Li Wang. A foundation model for enhancing magnetic resonance images and downstream segmentation, registration and diagnostic tasks. *Nature biomedical engineering*, 2024. URL <https://api.semanticscholar.org/CorpusID:274527575>.
- Junshen Xu, S. Lala, Borjan A. Gagoski, Esra Abaci Turk, Patricia Ellen Grant, Polina Golland, and Elfar Adalsteinsson. Semi-supervised learning for fetal brain mri quality assessment with roi consistency. *Medical image computing and computer-assisted intervention : MICCAI ... International Conference on Medical Image Computing and Computer-Assisted Intervention*, 12266:386–395, 2020. URL <https://api.semanticscholar.org/CorpusID:219981237>.
- Junshen Xu, Daniel Moyer, Borjan A. Gagoski, Juan Eugenio Iglesias, P. Ellen Grant, Polina Golland, and Elfar Adalsteinsson. Nesvor: Implicit neural representation for slice-to-volume reconstruction in mri. *IEEE Transactions on Medical Imaging*, 42:1707–1719, 2023. URL <https://api.semanticscholar.org/CorpusID:255917104>.