



HHS Public Access

Author manuscript

Med Image Comput Comput Assist Interv. Author manuscript; available in PMC 2023 April 25.

Published in final edited form as:

Med Image Comput Comput Assist Interv. 2022 September ; 13436: 3–13.

doi:10.1007/978-3-031-16446-0_1.

SVoRT: Iterative Transformer for Slice-to-Volume Registration in Fetal Brain MRI

Junshen Xu¹, Daniel Moyer², P. Ellen Grant^{3,4}, Polina Golland^{1,2}, Juan Eugenio Iglesias^{2,4,5,6}, Elfar Adalsteinsson^{1,7}

¹Department of Electrical Engineering and Computer Science, MIT, Cambridge, MA, USA

²Computer Science and Artificial Intelligence Laboratory, MIT, Cambridge, MA, USA

³Fetal-Neonatal Neuroimaging and Developmental Science Center, Boston Children's Hospital, Boston, MA, USA

⁴Harvard Medical School, Boston, MA, USA

⁵Centre for Medical Image Computing, Department of Medical Physics and Biomedical Engineering, University College London, UK

⁶Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital and Harvard Medical School, Boston, USA

⁷Institute for Medical Engineering and Science, MIT, Cambridge, MA, USA

Abstract

Volumetric reconstruction of fetal brains from multiple stacks of MR slices, acquired in the presence of almost unpredictable and often severe subject motion, is a challenging task that is highly sensitive to the initialization of slice-to-volume transformations. We propose a novel slice-to-volume registration method using Transformers trained on synthetically transformed data, which model multiple stacks of MR slices as a sequence. With the attention mechanism, our model automatically detects the relevance between slices and predicts the transformation of one slice using information from other slices. We also estimate the underlying 3D volume to assist slice-to-volume registration and update the volume and transformations alternately to improve accuracy. Results on synthetic data show that our method achieves lower registration error and better reconstruction quality compared with existing state-of-the-art methods. Experiments with real-world MRI data are also performed to demonstrate the ability of the proposed model to improve the quality of 3D reconstruction under severe fetal motion.

Keywords

Slice-to-volume Registration; 3D reconstruction; Fetal MRI; Transformer

1 Introduction

Volumetric reconstruction of fetal brain from multiple motion-corrupted stacks of MR slices is an important tool in studying fetal brain development [8,2,27]. Due to rapid and random fetal motion, fetal MRI are limited to fast acquisition techniques, such as the single-shot T2 weighted (SST2W) imaging that freezes in-plane motion. Even with such fast 2D sequences, fetal MRI is still vulnerable to inter-slice motion artifacts [19], leading to misalignment of slices in a stack. Moreover, the delay between slices due to safety constraints on specific absorption rate (SAR) [17] further worsen the situation. Therefore, Slice-to-Volume Registration (SVR) prior to 3D reconstruction of fetal brain is necessary. Manual SVR is usually infeasible in practice due to the magnitude of data involved. Although optimization-based SVR methods have successfully applied to 3D reconstruction of fetal brain [14,18,7], coarse alignment of slices is required to initialize the algorithm and the quality of reconstructed volume is highly dependent on the initial alignment. Hence, an automatic and accurate method for estimating slice transformations is crucial to 3D reconstruction of fetal brain.

In an attempt to speed up SVR of fetal MRI and improve its capture range, deep learning methods [11,23] have been proposed to predict rigid transformations of MR slices using Convolution Neural Networks (CNNs), which share similarity with camera pose prediction in computer vision [12,15]. Pei *et al.* [22] proposed a multi-task network to exploit semantic information in fetal brain anatomy which, however, requires annotations of segmentation maps. Moreover, these approaches process each slice independently, ignoring the dependencies between slices. Singh *et al.* [25] proposed a recurrent network to predict inter-slice motion in fetal MRI. In SVR of fetal ultrasound, Yeung *et al.* [31] tried to predict the 3D location of multiple slices simultaneously with an attention CNN.

Recently, Transformer models [28] and their variants have achieved astounding results in various fields [4,3]. The concept behind Transformers is to dynamically highlight the relevant features in input sequences with the self-attention mechanism, which demonstrates great capability of modeling long-distance dependencies and capturing global context. In SVR of fetal MRI, multiple stacks of slices are provided as inputs, which can also be modeled as a sequence of images. Multi-view information from stacks of slices with different orientations can be processed jointly to assist the SVR task.

Here, we propose a Slice-to-Volume Registration Transformer (SVoRT) to map multiple stacks of fetal MR slices into a canonical 3D space and to further initialize SVR and 3D reconstruction. As such, we present the following contributions: 1) We propose a Transformer-based network that models multiple stacks of slices acquired in one scan as a sequence of images and predicts rigid transformations of all the slices simultaneously by sharing information across the slices. 2) The model also estimates the underlying 3D volume to provide context for localizing slices in 3D space. 3) In the proposed model, slice transformations are updated in an iterative manner to progressively improve accuracy.

2 Methods

Given n acquired slices of a scan, $y = [y_1, \dots, y_n]$, the goal of SVoRT is to estimate the corresponding transformations $T = [T_1, \dots, T_n]$, *i.e.*, rotations and translations of acquisition planes, in a 3D canonical atlas space. However, unlike SVR problems [5,9] where a 3D volume exists as a reference for matching 2D slices, high quality 3D references are usually unavailable in SVR of fetal MR due to fetal motion. Therefore, instead of predicting the transformations alone, we also estimate the underlying volume x from the input slices, so that the estimated volume \hat{x} can provide 3D context to improve the accuracy of predicted transformations. In SVoRT, the estimated transformation \hat{T} and the estimated volume \hat{x} are updated alternately, generating a series of estimates, $(\hat{T}^0, \hat{x}^0), \dots, (\hat{T}^K, \hat{x}^K)$, where \hat{T}^0 and \hat{x}^0 are the initial guesses, and K is the number of iterations. The estimated transformations of the last iteration \hat{T}^K is used as the output of the model. Fig. 1 (a) shows the k -th iteration of SVoRT, which consists of two steps: 1) computing the new transformation \hat{T}^k given \hat{T}^{k-1} and \hat{x}^{k-1} from the previous iteration, and 2) estimating volume \hat{x}^k based on the new transformation \hat{T}^k .

2.1 Transformation Update

At the k -th iteration, the transformations are updated by $\hat{T}^k = \hat{T}^{k-1} + \Delta\hat{T}^k$. We propose a submodule named Slice-Volume Transformer (SVT) to regress the residual $\Delta\hat{T}^k$ given the set of input slice and the estimates from the previous iteration, $\Delta\hat{T}^k = \text{SVT}_T^k(y, \hat{T}^{k-1}, \hat{x}^{k-1})$. SVT, whose architecture is shown in Fig. 1 (b), aims to jointly extract features from stacks of slices and a 3D volume.

To relate the volume \hat{x}^{k-1} to the set of slices y with estimated transformations \hat{T}^{k-1} , SVT simulates slices from the volume following the forward model, $\hat{y}_i = D\hat{B}\hat{T}_i^{k-1}\hat{x}^{k-1}$, $i = 1, \dots, n$, where D and B are the operators for slice sampling and Point-Spread-Function (PSF) blurring respectively. The simulated slices \hat{y} provide views of the estimated volume \hat{x} at the estimated slice locations \hat{T}^{k-1} . The difference between \hat{y} and the original slices y can be used as a proxy indicator of registration accuracy and guide models to update the estimated transformations. To this end, we concatenate \hat{y} and y , and use a ResNet [10] to extract features $X^{\text{slice}} \in \mathbb{R}^{n \times d}$ from slices, where d is the number of features.

In addition to the image content, information about the position of the slice in the sequence and the estimated location in 3D space is injected for Transformers to encode spatial correlation of the input sequence, *e.g.*, adjacent slices in the same stacks are usually highly correlated, while stacks with different orientations provide complementary information. Each slice in the input sequence is associated with two indices, the index of the stack that the slice belongs to, and the index of the slice in the stack. Positional embeddings $X^{\text{pos}} \in \mathbb{R}^{n \times d}$ are generated from the current estimated transformation \hat{T}^{k-1} , the stack indexes, and the slice indexes using sinusoidal functions [28].

The slice features and the corresponding positional embeddings are added and provided to a Transformer with four encoders [28]. Each Transformer encoder consists of a multi-head attention module, a feed forward network, and two layer normalization [1]. Let $X = X^{\text{slices}} + X^{\text{pos}}$ be the input matrix. The multi-head attention module first projects X into three different spaces, $Q_j = XW_j^o$, $K_j = XW_j^k$, $V_j = XW_j^v$, $j = 1, \dots, h$, where weights $W_j^o, W_j^k, W_j^v \in \mathbb{R}^{d \times (d/h)}$ and h is the number of heads. Then, each head computes the output as $Y_j = \text{softmax}(Q_j K_j^T / \sqrt{d}) V_j$, where the softmax function is applied to each row. The outputs from all heads are aggregated and projected with a matrix $W^O \in \mathbb{R}^{d \times d}$, i.e., $Y = [Y_1, \dots, Y_h] W^O$. The feed forward network is a fully connected network used to extract deeper features. At the end, a linear layer is applied to regress the residual transformations from the output of Transformer.

2.2 Volume Estimation

The next step is to compute the new estimate of volume \hat{x}^k based on the updated transformations \hat{T}^k . One of the available methods is the PSF reconstruction [14], which aligns the slices in 3D space based on transformations \hat{T}^k and interpolates the volume with the PSF kernel. However, there are two disadvantages of this approach. First, it over smooths the reconstructed volume and leads to a loss of image detail. Second, it fails to exclude slices with large transformation error during reconstruction, resulting in artifacts in the reconstructed volume.

To address these problems, we use another SVT in the volume estimation step to predict weights of slices, $w = [w_1, \dots, w_n]$, where $w_j \in [0, 1]$ represents the image quality of the i -th slice. The SVT here shares the same architecture as the one in the transformation update step, but has different inputs. Specifically, in the inputs to SVT, \hat{T}^{k-1} and \hat{x}^{k-1} are replaced with the updated transformations \hat{T}^k and the PSF reconstruction result \hat{x}_{PSF} respectively, $w^k = \text{SVT}_x^k(\hat{T}^k, \hat{x}_{\text{PSF}}^k, y)$, where we denote the SVT in volume estimation as SVT_x^k . To compute the new estimated volume \hat{x}^k , we solve an inverse problem,

$$\hat{x}^k = \underset{x}{\text{argmin}} \sum_{i=1}^n w_i^k \|DB\hat{T}_i^k x - y_i\|_2^2, \quad (1)$$

where the weights help exclude outliers during volume estimation. Since the closed form solution involves inverting a very large matrix, we instead employ a conjugate gradient (CG) method to compute \hat{x}^k . Note that all operations (matrix multiplication, addition and scalar division) in CG are differentiable, so the gradient with respect to w can be computed via automatic differentiation.

2.3 Training

Data: Supervised learning of SVoRT requires the ground truth transformation of each slice. However, annotating the 3D location of a 2D MR slice is very challenging. Instead, we artificially sample 2D slices as training data from high quality MR volumes of fetal brain

reconstructed from data with little fetal motion. The orientations of stack, *i.e.*, the normal vector and in-plane rotation, are randomly sampled as in [11] so that the dataset captures a wide range of rigid transformations. To Bridge the gap between the synthetic data and real MR scans and improve the generalization capability of networks, we also adopt various data augmentation and MR artifact simulation techniques [13], including image noise, bias field, signal void artifacts, and contrast jitter.

Representations of transformations: Various representations are available for the describing the location of a plane in 3D space. For example, the Euler angles, the affine matrix, or the Cartesian coordinates of 3 points, called anchor points, within the plane. Previous works have demonstrated that the anchor point representation yields the highest accuracy for deep learning based SVR methods. Following [11], we use the center, the bottom right and left corners of a plane as the anchor points to define the rigid transformation.

Loss functions: During training, we apply the L2 loss between the predicted and target anchor points for transformation prediction,

$$\mathcal{L}_T^k = \|\hat{P}_1^k - P_1\|_2^2 + \|\hat{P}_2^k - P_2\|_2^2 + \|\hat{P}_3^k - P_3\|_2^2,$$

where $\hat{P}_1^k, \hat{P}_2^k, \hat{P}_3^k$ are the predicted coordinates of the three anchor points in the k -th iteration, and P_1, P_2, P_3 are the ground truth locations. As for volume estimation, the L1 loss between the k -th estimated volume and the target volume is used, $\mathcal{L}_x^k = \|x^k - x\|_1$. The total loss \mathcal{L} is the sum of the losses in all iterations, $\mathcal{L} = \sum_{k=1}^K \mathcal{L}_T^k + \lambda \sum_{k=1}^K \mathcal{L}_x^k$, where λ is a weight determining the relative contribution of the L1 loss.

3 Experiments and Results

3.1 Experiment Setup

We evaluate the models on the FeTA dataset [21], which consists of 80 T2-weighted fetal brain volumes with gestational age between 20 and 35 weeks. The dataset is split into training (68 volumes) and test (12 volumes) sets. The volumes are registered to a fetal brain atlas [8], and resampled to 0.8 mm isotropic. We simulate 2D slices with resolution of 1 mm \times 1 mm, slice thickness between 2.5 and 3.5 mm, and size of 128 \times 128. Each training sample consists of 3 image stacks in random orientations and each stack has 15 to 30 slices. Fetal brain motion is simulated as in [30] to perturb the transformations of slices. In the process of training, random samples are generated on the fly, while for testing, 4 different samples are generated for each test subject, resulting in 48 test cases.

To demonstrate SVoRT can generalize well to real-world data and help initialize SVR for cases with severe fetal motion, we test the trained models with data acquired in real fetal MRI scans. Scans were conducted at Boston Children’s Hospital. MRI data were acquired using the HASTE sequence [6] with slice thickness of 2 mm, resolution of 1 mm \times 1 mm, size of 256 \times 256, TE = 119 ms, and TR = 1.6 s. The real MR dataset not only has different contrast, but also undergoes more realistic artifacts and fetal motion compared

to the synthetic data. In the experiments, SVRnet [11] and PlaneInVol [31] are used as baselines. SVRnet predicts the transformation of each slice independently with a VGGNet [24], while PlaneInVol uses an attention CNN which compares pairs of slices, and learns to map a set of slices to 3D space. For SVoRT, we set the initial estimates \hat{T}^0 and \hat{x}^0 to identity transformation and zero respectively and set $\lambda = 10^3$ and $K = 3$. All neural networks are implemented with PyTorch [20] and trained on a Nvidia Tesla V100 GPU with an Adam optimizer [16] for 2×10^5 iterations. We used an initial learning rate of 2×10^{-4} which linearly decayed to zero. The reference implementation for SVoRT is available on GitHub⁸.

3.2 Simulated Data

To evaluate the accuracy of the predicted transformation for different models, we use the Euclidean Distance (ED) of anchor points, and the Geodesic Distance (GD) in $SO(3)$: $ED = \frac{1}{3} \sum_{j=1}^3 \|\hat{P}_j - P_j\|_2$ and $GD = \arccos\left(\frac{Tr(R) - 1}{2}\right)$, where R is the rotation matrix from the predicted plane to the target, representing the rotation error. We also extract slices from the ground truth volume x according to the estimated transformations \hat{T} and compare them to the original slices y . Comparison is performed via Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity (SSIM) [29]. To further examine the model, we use the predicted transformations to initialize a 3D reconstruction algorithm [14] and compute PSNR and SSIM between the reconstructed volumes and the targets.

Table 1 reports the mean and standard deviations of quantitative metrics for different models on the test set of the simulated data. Our proposed method outperforms both SVRnet and PlaneInVol, which only leverage the intensity information of slices. As shown in Fig. 2 (a), the transformation errors for SVRnet and PlaneInVol increase with the distance to the center of 3D space, since the slices near the boundary of fetal brain contain little content and can be ambiguous. However, by exploiting the positional information of slice in the input sequence, SVoRT is able to register such cases better and lead to lower errors.

Ablation studies are also performed by removing the positional embedding (w/o PE) and the volume estimation (w/o Vol), and using fewer iterations ($K = 1, 2$) in SVoRT. Results indicate that the positional embedding serves as a prior for the relative locations of slices in a stack, which facilitates the registration process. The auxiliary volume estimation improves the accuracy of transformation prediction by providing 3D context. Moreover, the iterative update enable the model to progressively refine the predictions. We test SVoRT with different numbers of input stacks (Fig. 2 (b)). With more input stacks, SVoRT receives more different views of the 3D volume, and therefore achieves lower registration error. We also compare different volume estimation methods: i) the proposed estimate \hat{x} , ii) the solution to the inverse problem in Eq. (1) with equal weight, $\hat{x}(w = 1)$, and iii) the PSF reconstruction, \hat{x}_{PSF} . As shown in Fig. 2 (c), the proposed method achieves the highest PSNR. Fig. 2 (d) visualizes an example attention matrix generated by the last Transformer encoder. The 3D reconstruction results in Fig. 3 show that SVoRT also yields better perceptual quality

⁸ <https://github.com/daviddmc/SVoRT>

compared with other state-of-the-art methods, in consequence of fewer slice misalignment in the initialization of SVR.

3.3 Real Fetal MR Data

We collect 3 orthogonal motion-corrupted stacks of MR slices from two subjects respectively. For preprocessing, bias fields are corrected [26], and the brain ROI is manually segmented from each slice. SVRnet, PlaneInVol and SVoRT are used to predict the transformations and initialize the SVR algorithm. For comparison, we also apply the SVR algorithm to the input stacks directly without deep learning based initialization (SVR only). Note that the results of “SVR only” are reconstructed in the subject space. We further register them to the atlas space for visualization. Fig. 4 shows that volumes reconstructed by SVR alone suffer from severe image artifacts due to slice misalignment caused by drastic fetal motion. SVRnet and PlaneInVol are incapable of generalizing to real MR data and fail to provide a useful initialization for SVR. In comparison, the estimated transformations of SVoRT are more accurate and the corresponding reconstructed volume presents better perceptual quality. Results indicate that SVoRT learns more robust features from synthetic data, and therefore generalizes well in the presence of real-world noise and artifacts. Moreover, the average inference time of SVoRT for each subject is 0.8 s, which is negligible compared with SVR algorithms that usually take minutes even on GPUs. SVoRT potentially enables high quality 3D reconstruction of fetal MRI in the case of severe fetal motion.

4 Conclusion

In this work, we propose a novel method for slice-to-volume registration in fetal brain MRI using Transformers. By jointly processing the stacks of slices as a sequence, SVoRT registers each slice by utilizing context from other slices, resulting in lower registration error and better reconstruction quality. In addition, we introduce an auxiliary task of volume estimation and update the transformation iteratively to improve registration accuracy. Evaluations show that SVoRT learns more robust features so that, by training on simulated data, it generalizes well to data acquired in real scans. SVoRT provides a robust and accurate solution to the initialization of 3D fetal brain reconstruction.

Acknowledgements

This research was supported by NIH U01HD087211, R01EB01733, HD100009, NIBIB R01EB032708, NIBIB NAC P41EB015902, R01AG070988, RF1MH123195, ERC Starting Grant 677697, ARUK-IRG2019A-003.

References

1. Ba JL, Kiros JR, Hinton GE: Layer normalization. arXiv preprint arXiv:1607.06450 (2016)
2. Benkarim O, Piella G, Rekik I, Hahner N, Eixarch E, Shen D, Li G, Ballester MAG, Sanroma G: A novel approach to multiple anatomical shape analysis: application to fetal ventriculomegaly. *Medical image analysis* 64, 101750 (2020) [PubMed: 32559594]
3. Devlin J, Chang MW, Lee K, Toutanova K: Bert: Pre-training of deep bidirectional transformers for language understanding (2019)
4. Dosovitskiy A, Beyer L, Kolesnikov A, Weissenborn D, Zhai X, Unterthiner T, Dehghani M, Minderer M, Heigold G, Gelly S, et al. : An image is worth 16×16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)

5. Esteban J, Grimm M, Unberath M, Zahnd G, Navab N: Towards fully automatic X-ray to CT registration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 631–639. Springer (2019)
6. Gagoski B, Xu J, Wighton P, Tisdall MD, Frost R, Lo WC, Golland P, van Der Kouwe A, Adalsteinsson E, Grant PE: Automated detection and reacquisition of motion-degraded images in fetal haste imaging at 3 t. *Magnetic Resonance in Medicine* (2021)
7. Gholipour A, Estroff JA, Warfield SK: Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain MRI. *IEEE transactions on medical imaging* 29(10), 1739–1758 (2010) [PubMed: 20529730]
8. Gholipour A, Rollins CK, Velasco-Annis C, Oualam A, Akhondi-Asl A, Afacan O, Ortinau CM, Clancy S, Limperopoulos C, Yang E, et al. : A normative spatiotemporal MRI atlas of the fetal brain for automatic segmentation and analysis of early brain growth. *Scientific reports* 7(1), 1–13 (2017) [PubMed: 28127051]
9. Gillies DJ, Gardi L, De Silva T, Zhao S.r., Fenster A: Real-time registration of 3d to 2d ultrasound images for image-guided prostate biopsy. *Medical physics* 44(9), 4708–4723 (2017) [PubMed: 28666058]
10. He K, Zhang X, Ren S, Sun J: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
11. Hou B, Khanal B, Alansary A, McDonagh S, Davidson A, Rutherford M, Hajnal JV, Rueckert D, Glocker B, Kainz B: 3-d reconstruction in canonical co-ordinate space from arbitrarily oriented 2-d images. *IEEE transactions on medical imaging* 37(8), 1737–1750 (2018) [PubMed: 29994453]
12. Hou B, Miolane N, Khanal B, Lee MC, Alansary A, McDonagh S, Hajnal JV, Rueckert D, Glocker B, Kainz B: Computing cnn loss and gradients for pose estimation with riemannian geometry. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 756–764. Springer (2018)
13. Iglesias JE, Billot B, Balbastre Y, Tabari A, Conklin J, González RG, Alexander DC, Golland P, et al. : Joint super-resolution and synthesis of 1 mm isotropic mp-rage volumes from clinical MRI exams with scans of different orientation, resolution and contrast. *NeuroImage* 237, 118206 (2021) [PubMed: 34048902]
14. Kainz B, Steinberger M, Wein W, Kuklisova-Murgasova M, Malamateniou C, Keraudren K, Torsney-Weir T, Rutherford M, Aljabar P, Hajnal JV, et al. : Fast volume reconstruction from motion corrupted stacks of 2d slices. *IEEE transactions on medical imaging* 34(9), 1901–1913 (2015) [PubMed: 25807565]
15. Kendall A, Grimes M, Cipolla R: Posenet: A convolutional network for real-time 6-dof camera relocalization. In: Proceedings of the IEEE international conference on computer vision. pp. 2938–2946 (2015)
16. Kingma DP, Ba J: Adam: A method for stochastic optimization (2017)
17. Krishnamurthy U, Neelavalli J, Mody S, Yeo L, Jella PK, Saleem S, Korzeniewski SJ, Cabrera MD, Ehterami S, Bahado-Singh RO, et al. : Mr imaging of the fetal brain at 1.5 t and 3.0 t field strengths: comparing specific absorption rate (sar) and image quality. *Journal of perinatal medicine* 43(2), 209–220 (2015) [PubMed: 25324440]
18. Kuklisova-Murgasova M, Quaghebeur G, Rutherford MA, Hajnal JV, Schnabel JA: Reconstruction of fetal brain MRI with intensity matching and complete outlier removal. *Medical image analysis* 16(8), 1550–1564 (2012) [PubMed: 22939612]
19. Malamateniou C, Malik S, Counsell S, Allsop J, McGuinness A, Hayat T, Broadhouse K, Nunes R, Ederies A, Hajnal J, et al. : Motion-compensation techniques in neonatal and fetal mr imaging. *American Journal of Neuroradiology* 34(6), 1124–1136 (2013) [PubMed: 22576885]
20. Paszke A, Gross S, Chintala S, Chanan G, Yang E, DeVito Z, Lin Z, Desmaison A, Antiga L, Lerer A: Automatic differentiation in pytorch (2017)
21. Payette K, de Dumast P, Kebiri H, Ezhov I, Paetzold JC, Shit S, Iqbal A, et al. : An automatic multi-tissue human fetal brain segmentation benchmark using the fetal tissue annotation dataset. *Scientific Data* 8(1), 1–14 (2021) [PubMed: 33414438]

22. Pei Y, Wang L, Zhao F, Zhong T, Liao L, Shen D, Li G: Anatomy-guided convolutional neural network for motion correction in fetal brain MRI. In: International Workshop on Machine Learning in Medical Imaging. pp. 384–393. Springer (2020)
23. Salehi SSM, Khan S, Erdogmus D, Gholipour A: Real-time deep pose estimation with geodesic loss for image-to-template rigid registration. IEEE transactions on medical imaging 38(2), 470–481 (2018) [PubMed: 30138909]
24. Simonyan K, Zisserman A: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
25. Singh A, Salehi SSM, Gholipour A: Deep predictive motion tracking in magnetic resonance imaging: application to fetal imaging. IEEE Transactions on Medical Imaging 39(11), 3523–3534 (2020) [PubMed: 32746102]
26. Tustison NJ, Avants BB, Cook PA, Zheng Y, Egan A, Yushkevich PA, Gee JC: N4itk: improved n3 bias correction. IEEE transactions on medical imaging 29(6), 1310–1320 (2010) [PubMed: 20378467]
27. Vasung L, Turk EA, Ferradal SL, Sutin J, Stout JN, Ahtam B, Lin PY, Grant PE: Exploring early human brain development with structural and physiological neuroimaging. Neuroimage 187, 226–254 (2019) [PubMed: 30041061]
28. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I: Attention is all you need. Advances in neural information processing systems 30 (2017)
29. Wang Z, Bovik AC, Sheikh HR, Simoncelli EP: Image quality assessment: from error visibility to structural similarity. IEEE transactions on image processing 13(4), 600–612 (2004) [PubMed: 15376593]
30. Xu J, Abaci Turk E, Grant PE, Golland P, Adalsteinsson E: Stress: Super-resolution for dynamic fetal MRI using self-supervised learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 197–206. Springer (2021)
31. Yeung PH, Aliasi M, Papageorghiou AT, Haak M, Xie W, Namburete AI: Learning to map 2d ultrasound images into 3d space with minimal human annotation. Medical Image Analysis 70, 101998 (2021) [PubMed: 33711741]

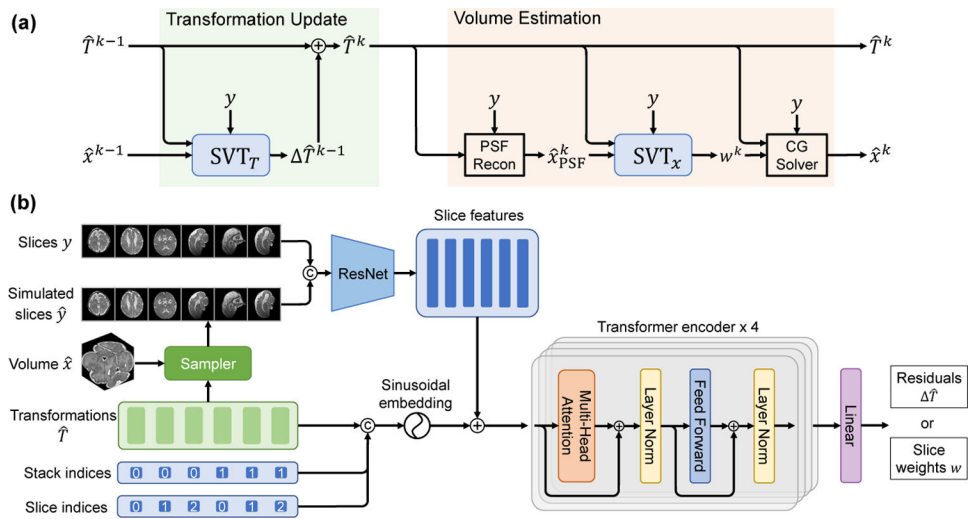


Fig. 1. (a) The k -th iteration of SVoRT. (b) The architecture of SVT.




Fig. 2.

(a) Medians of ED for slices sampled at different locations at the atlas space. Error bands indicate 25 and 75 percentiles. (b) Mean values of ED for SVoRT models with different numbers of input stacks. Error bars indicate standard deviations (c) PSNRs of different volume estimation methods. (d) An example heatmap of the self-attention weight matrix averaged over all the heads.




Fig. 3. Example reconstructed volumes and reference volumes of the test set.




Fig. 4. Reconstructed volumes for different methods of real MR data.

Table 1.

Mean values of quantitative metrics for different models (standard deviation in parentheses). ↓ indicates lower values being more accurate, vice versa. The best results are highlighted in bold.

Method	Transformation		Slice		Volume	
	ED(mm)↓	GD(rad)↓	PSNR↑	SSIM↑	PSNR↑	SSIM↑
SVRnet	12.82 (5.69)	.256 (.150)	20.53 (1.62)	.823 (.061)	19.54 (1.52)	.669 (.116)
PlaneInVol	12.49 (6.73)	.244 (.213)	19.96 (1.73)	.808 (.069)	18.98 (1.62)	.615 (.139)
SVoRT	4.35 (0.90)	.074 (.017)	25.26 (1.86)	.916 (.034)	23.32 (1.42)	.858 (.037)
w/o PE	9.97 (6.28)	.194 (.179)	21.44 (2.08)	.841 (.064)	20.74 (1.49)	.742 (.096)
w/o Vol	5.09 (1.05)	.088 (.020)	23.97 (1.68)	.894 (.040)	22.89 (1.37)	.844 (.043)
$K = 1$	5.99 (1.16)	.103 (.024)	23.02 (1.67)	.876 (.047)	22.57 (1.21)	.836 (.041)
$K = 2$	5.65 (1.07)	.097 (.022)	23.25 (1.84)	.878 (.048)	22.64 (1.50)	.837 (.043)