

Fetuses Made Simple: Modeling and Tracking of Fetal Shape and Pose

Yingcheng Liu¹, Peiqi Wang¹, Sebastian Diaz¹, Esra Abaci Turk², Benjamin Billot³, P. Ellen Grant², and Polina Golland¹

¹ Computer Science and Artificial Intelligence Lab, MIT, Cambridge, USA

² Boston Children’s Hospital and Harvard Medical School, Boston, USA

³ Inria, Epione team, Sophia-Antipolis, France

liuyc@mit.edu

Abstract. Analyzing fetal body motion and shape is paramount in prenatal diagnostics and monitoring. Existing methods for fetal MRI analysis mainly rely on anatomical keypoints or volumetric body segmentations. Keypoints simplify body structure to facilitate motion analysis, but may ignore important details of full-body shape. Body segmentations capture complete shape information but complicate temporal analysis due to large non-local fetal movements. To address these limitations, we construct a 3D articulated statistical fetal body model based on the Skinned Multi-Person Linear Model (SMPL). Our algorithm iteratively estimates body pose in the image space and body shape in the canonical pose space. This approach improves robustness to MRI motion artifacts and intensity distortions, and reduces the impact of incomplete surface observations due to challenging fetal poses. We train our model on segmentations and keypoints derived from 19,816 MRI volumes across 53 subjects. Our model captures body shape and motion across time series and provides intuitive visualization. Furthermore, it enables automated anthropometric measurements traditionally difficult to obtain from segmentations and keypoints. When tested on unseen fetal body shapes, our method yields a surface alignment error of 3.2 mm for 3 mm MRI voxel size. To our knowledge, this represents the first 3D articulated statistical fetal body model, paving the way for enhanced fetal motion and shape analysis in prenatal diagnostics. The code is available at <https://github.com/MedicalVisionGroup/fetal-smpl>.

Keywords: Fetal MRI · Pose estimation · Shape estimation · 3D body model · SMPL (Skinned Multi-Person Linear Model)

1 Introduction

Analyzing fetal body motion and shape is crucial in prenatal diagnosis and monitoring. For example, fetal movement is related to the function of the fetal central nervous system [19], and anatomical growth patterns provide key biomarkers of fetal health [17]. In the current clinical workflow, most assessments are performed visually with 2D ultrasound. When ultrasound findings are inconclusive,




Fig. 1. Fetal rave. We learn a 3D articulated statistical fetal body shape model from 20k volumes of MRI time series in 53 subjects. (A): 3D MRI time series with fetal body segmentations. (B): Model aligned to volumetric frames in A. (C): Overview of our model aligned to all volumes in our dataset.

3D MRI is an established solution that visualizes fetal anatomy fully in 3D. Recent advances in machine learning have enabled automatic fetal MRI analysis [19,22,20,3,16]. These methods typically model the fetal body as a set of anatomical landmarks [19,22,20] or rigid body parts [3,16]. Keypoint-based approaches facilitate motion analysis by representing the body as a sparse “stick figure,” but this simplification may ignore important full-body shape information. Alternatively, volumetric segmentation captures the full body shape [18,15]. However, motion analysis with segmentation is challenging because the fetal body undergoes large deformations. Existing deformable registration methods often assume small local deformations and struggle to handle the wide variability in pose and large non-local deformations of the fetal body [6].

In this paper, we aim to overcome these limitations by building a surface-based kinematic fetal body model using segmentation and keypoints from MRI time series. Our model is based on Skinned Multi-Person Linear Model (SMPL) [11]. We represent the fetal body as a collection of connected rigid parts parameterized by a kinematic tree. The kinematic tree enables us to align fetal body into a standardized pose, where we define a canonical space and subject-specific body shape. This approach decouples the representations of the pose-induced deformation and intrinsic body shape. Once trained and aligned to images, our model facilitates the extraction of surface correspondence between frames for body motion analysis.

Building a fetal SMPL model using MRI is challenging. In adult SMPL model training, the correspondence between the model and data surface is known [11]. When the correspondence is not available, existing methods often require the subject to lie in a standard pose [1]. The body shape is then estimated by minimizing the Chamfer distance between the model and data surfaces [1]. In our application, placing the fetuses in a shared standard pose during the MRI imaging procedure is impossible. Furthermore, our images are low-resolution and frequently contain noise from motion artifacts or maternal breathing. Surface

observations are also often incomplete. Moreover, body parts (*e.g.*, legs and abdomen) frequently touch, making it difficult to extract the skin surface from volumetric segmentations.

To tackle these challenges, we iteratively estimate body shape and pose using a coordinate descent algorithm. First, we estimate pose using an initial body shape and “unpose” the observed body surfaces in all frames into the common canonical space. We then reconstruct the subject-specific canonical shape using all the unposed body surfaces. This enables us to aggregate shape information across frames and mitigate the impact of noisy, incomplete surface observations in individual frames. Subsequently, the pose estimate is refined using the updated shape estimates. We repeat these steps until shape and pose estimates converge. This approach is inspired by prior work on adult and newborn body modeling [1,7,8], adapted specifically for fetal MRI. Additionally, we improve the method by deriving exact coordinate transformations with inverse kinematics, which supports accurate mapping of the segmentation vertices to the canonical space.

We train fetal SMPL using segmentations and keypoints derived from 19,816 volumes of Echo Planar Imaging (EPI) MRI data across 53 subjects. Fig. 1 illustrates the dynamics of our model and an overview of the registered model. The model captures diverse fetal shape and pose, providing new tools for fetal MRI analysis. Experiments show that our model explains unseen shapes with a surface alignment error of 3.2 mm in volumes with 3 mm isotropic voxels. The resulting anthropometric measurements illustrate potential applications for developmental abnormality detection.

2 Method

2.1 SMPL Background

The SMPL model [11] is a 3D surface-based human body model with $N = 6,890$ vertices. The model is parameterized by a kinematic tree containing $K = 23$ joints. The body pose $\theta \in \mathbb{R}^{72}$ is represented as rotations at these K joints in angle-axis format, plus a global rotation. The global translation is represented by $t \in \mathbb{R}^3$. The canonical pose θ^* (T-pose) standardizes body poses across subjects. PCA is applied to a dataset of T-pose shapes, yielding a mean shape $\bar{\mathbf{T}}$ and shape principal components (PCs) $B_S(\beta)$ that capture shape variation over a population. Subject-specific canonical shape $\mathbf{T}_S = \bar{\mathbf{T}} + B_S(\beta) \in \mathbb{R}^{3N}$ is computed from a shape vector $\beta \in \mathbb{R}^{10}$. Joint locations $\mathbf{J} = \mathcal{J}\mathbf{T}_S$ are linearly regressed from the canonical shape using a subject-agnostic joint regression matrix $\mathcal{J} \in \mathbb{R}^{3K \times 3N}$. In addition, the model defines pose blend shapes $B_P(\theta; \mathcal{P}) = (R(\theta) - R(\theta^*))\mathcal{P}$ to handle pose-induced deformations, where $R \in \mathbb{R}^{9K}$ are the vectorized rotation matrices of all joints, and $\mathcal{P} \in \mathbb{R}^{3N \times 9K}$ represents pose-induced deformations learned from data. $T_P(\beta, \theta) = \bar{\mathbf{T}} + B_S(\beta) + B_P(\theta)$ includes shape- and pose-induced corrections. Each vertex $\mathbf{v}_n \in T_P(\beta, \theta)$ is transformed using linear blend skinning $\mathbf{v}'_n = \sum_{k=1}^K w_{k,n} G'_k \mathbf{v}_n$, where $w_{k,n}$ is the weights of joint k for

vertex n , and G'_k is the relative transformation of joint k . We refer the readers to the SMPL paper for more details [11].

2.2 Training Fetal Model

For subject i and volume j in a given time series, our training data comprises keypoints $\mathbf{K}_{i,j} \in \mathbb{R}^{3 \times 16}$ (which do not necessarily map one-to-one to SMPL joints) and body meshes $\mathbf{V}_{i,j} \in \mathbb{R}^{3 \times N_j}$. We use the marching cubes algorithm [12] to extract surface vertices from the volumetric segmentations. To construct the fetal body model, we estimate the canonical shape $\mathbf{T}_S$ for each subject and the body pose θ for each frame in the time series. Applying PCA to $\mathbf{T}_S$ across all subjects yields the mean shape $\bar{\mathbf{T}}$ and shape PCs B_S . We learn the pose blend shapes $\mathcal{P}$ and the anatomical keypoint regressor $\mathcal{J}_k$ from our fetal data.

We use a coordinate descent algorithm to alternate between the estimation of the body pose and canonical shape. We estimate the body pose using initial canonical shapes. Then, we “unpose” surface vertices to the canonical space using the most recent body pose estimate. Next, we refine canonical shape estimates using all unposed vertices across the entire time series. We iterate these steps until convergence. In the following sections, we describe each step in more detail.

Initialization. We start the iterative process by registering the SMIL model [7] (*i.e.*, an infant SMPL model) to segmentation vertices $\mathbf{V}_{i,j}$ and keypoints $\mathbf{K}_{i,j}$. We fix SMIL model’s $\{\mathbf{T}, \mathcal{S}, \mathcal{W}\}$ and optimize pose $\theta_{i,j}$, translation $t_{i,j}$, and shape parameters β_i for each subject i by minimizing

$$E_{\text{init}} = \sum_j (E_v + \lambda_k E_k) + \lambda_{\text{shape}} E_{\text{shape}} + \lambda_{\text{smooth}} E_{\text{smooth}} + \lambda_{\text{prior}} E_{\text{prior}}. \quad (1)$$

The vertex term E_v is the two-sided Chamfer distance between the model and segmentation vertices $\mathbf{V}_{i,j}$. The keypoint term $E_k = \sum_k \rho(\|\mathcal{J}_k \mathbf{T}_{i,j} - \mathbf{K}_{i,j}\|)$ uses the generalized robust loss function [2] $\rho(\cdot)$ to measure distances between the model-predicted keypoints and the annotated keypoints $\mathbf{K}_{i,j}$. Transformed model vertices are given by $\mathbf{T}_{i,j} = W(T_{S,i} + B_P(\theta_{i,j}), \mathbf{J}, \theta_{i,j}, \mathcal{W}) + t_{i,j}$. $E_{\text{shape}} = \|\beta_i\|^2$ regularizes shape vectors. We add a temporal smoothness term E_{smooth} that imposes ℓ_2 penalty on the differences in pose $\theta_{i,j}$ and translation $t_{i,j}$ between consecutive frames to encourage smooth motion. To ensure the pose remains kinematically feasible, we encourage the predicted poses to be close to a multivariate Gaussian distribution measured using the squared Mahalanobis distance $E_{\text{prior}} = \sum_j (\theta_{i,j} - \bar{\theta})^\top \Sigma^{-1} (\theta_{i,j} - \bar{\theta})$ where $\bar{\theta}$ denotes the mean pose and Σ represents the covariance matrix of the SMIL pose prior. We initialize the pose blend shapes $\mathcal{P}$ to zero and the anatomical keypoint regressor $\mathcal{J}_k$ to be equal to the corresponding joints from the adult regressor $\mathcal{J}_A$. λ_k , λ_{shape} , λ_{smooth} , and λ_{prior} are the corresponding regularization coefficients.

Step 1: Pose. We fix the subject-specific canonical shape $\mathbf{T}_{S,i}$, the pose blend shapes $\mathcal{P}$, and the anatomical keypoint regressor $\mathcal{J}_k$, and optimize the pose

parameters $\theta_{i,j}$ and translations $t_{i,j}$ for each subject i and frame j to minimize:

$$E_{\text{pose}} = E_v + \lambda_k E_k + \lambda_{\text{smooth}} E_{\text{smooth}} + \lambda_{\text{prior}} E_{\text{prior}}, \quad (2)$$

with all terms defined above in the initialization step. We initialize the pose parameters $\theta_{i,j}$ and translations $t_{i,j}$ using the results from the previous step.

Step 2: Pose Blend Shapes and Anatomical Keypoint Regressor. Next, we fix the pose parameters $\theta_{i,j}$ and the canonical shapes $\mathbf{T}_{S,i}$, and optimize $\mathcal{P}$ by minimizing:

$$E_{\mathcal{P}} = \sum_{i,j} E_v + \lambda_P E_P + \lambda_{\text{elast}} E_{\text{elast}}, \quad (3)$$

where $E_P = \|\mathcal{P}\|_F^2$ penalizes the quadratic Frobenius norm of the pose blend shapes to prevent overfitting. The elastic regularization term $E_{\text{elast}} = \sum_e \|(\mathbf{T}_{S,i})_e - (\mathbf{T}_{S,i}^0)_e\|^2$ penalizes changes in edge vectors between the optimized meshes and their initial value. Here, $(\mathbf{T}_{S,i})_e$ and $(\mathbf{T}_{S,i}^0)_e$ denote edge vectors of the optimized and initial meshes, respectively. λ_P and λ_{elast} are corresponding weights.

Next, we update the anatomical keypoint regressor $\mathcal{J}_k$. We fix the pose parameters $\theta_{i,j}$, the canonical shapes $\mathbf{T}_{S,i}$, and the pose blend shapes $\mathcal{P}$, and optimize $\mathcal{J}_k$ by minimizing:

$$E_{\text{key}} = \sum_{i,j} E_k + \lambda_J E_J, \quad (4)$$

where $E_J = \|\mathcal{J}_k - \mathcal{J}_A\|^2$ is a regularization term designed to keep the anatomical keypoint regressor $\mathcal{J}_k$ close to the adult model regressor $\mathcal{J}_A$. λ_J is the corresponding weight.

Step 3: Canonical Shape. Next, we estimate canonical shapes $\mathbf{T}_{S,i}$ using the unposed segmentation vertices $\mathbf{V}_{i,j}$. For each vertex $\mathbf{v} \in \mathbf{V}_{i,j}$, we find the nearest model vertex $\mathbf{u}_k \in \mathbf{T}_{i,j}$ and apply the inverse transformation to map it to the canonical space:

$$\tilde{\mathbf{v}} = (G'_k(\theta_{i,j}, \mathbf{J}))^{-1} (\mathbf{v} - t_{i,j}). \quad (5)$$

We subtract the corresponding pose blend shape $B_P(\theta_{i,j})$ from the unposed vertices to isolate the canonical shape contribution. Then, we optimize the canonical shape $\mathbf{T}_{S,i}$:

$$E_{\text{shape}} = \sum_i (E_{\text{Chamfer}} + \lambda_{\text{elast}} E_{\text{elast}}), \quad (6)$$

where E_{Chamfer} is the two-sided Chamfer distance between the canonical shape $\mathbf{T}_{S,i}$ and the unposed surface vertices. The elastic regularization term E_{elast} is defined in Step 2 and λ_{elast} is its weighting factor.

Mean Canonical Shape $\bar{\mathbf{T}}$ and Shape Principal Components $\mathcal{S}$. We repeat the above pose and shape estimation steps (Steps 1-3) until convergence. Finally, we compute the mean canonical shape $\bar{\mathbf{T}}$ by averaging the optimized $\mathbf{T}_{S,i}$ across all subjects. We then apply robust PCA [5] to the residuals to derive the shape blend shapes $\mathcal{S}$. The mean canonical shape $\bar{\mathbf{T}}$, shape blend shapes $\mathcal{S}$, pose blend shapes $\mathcal{P}$, and the anatomical landmark regressor $\mathcal{J}_k$ together constitute our fetal body model.

New Subjects. For a new subject, we register the fetal model to segmentation vertices $\mathbf{V}_{i,j}$ and keypoint time series $\mathbf{K}_{i,j}$. This process mimics the model training procedure but keeps the model parameters fixed. We initialize the pose $\theta_{i,j}$ and the shape β_i parameters following the “Initialization” step using the trained fetal model. We then unpose the vertices to the canonical space. Next, we optimize the shape parameters β_i using the two-sided Chamfer distance to estimate the canonical shape $\mathbf{T}_S$. Finally, we refine the pose by optimizing $\theta_{i,j}$ with updated shape $\mathbf{T}_S$. Using the canonical shape $\mathbf{T}_S$, we assess clinically relevant anthropometric parameters, including body length (BL), head circumference (HC), abdominal circumference (AC), body volume (BV), and femur length (FL).

Implementation Details. We use the Adam optimizer [10] with a learning rate of $1e^{-3}$. We adopt the PyTorch SMPL-X implementation of SMPL [13]. For efficient 3D nearest neighbor search, we use the KD-tree implementation in PyTorch3D [14]. We use AltViewer [9] to visualize the meshes and keypoints.

3 Experiments and Results

Data and Evaluation. Our dataset comprises two cohorts: a research cohort and a clinical cohort, both from Boston Children’s Hospital. The research cohort includes 53 singleton subjects with gestational age (GA) of 24 – 37 weeks, imaged with EPI time-series of 216 ± 125 volumetric frames. The clinical cohort consists of 11 subjects, including 6 diagnosed with Chiari II malformation and 5 non-Chiari controls. The median (IQR) of GA among the clinical cohort is 24wk (22-35wk). All scans were acquired on a 3T Siemens Skyra scanner.

To test the model’s generalization, we created the “New Shape” and “New Pose” test sets. We randomly selected 20% subjects from the research cohort to form the “New Shape” subset. For the remaining subjects, we leave out 10% of the frames in each time series to form the “New Pose” dataset. The “New Pose” dataset evaluates generalization to unseen poses, and the “New Shape” dataset assesses generalization to unseen shapes. We repeat the above data selection process to obtain five non-overlapping “New Shape” test sets. This ensures each subject contributes once to the “New Shape” dataset.

We curated a ground truth fetal body segmentation dataset and trained a UNet-based model to obtain fetal body segmentations [4]. Keypoints were obtained using a pre-trained model [21,22]. The trained fetal body model was




Fig. 2. Alignment error statistics. (A) Alignment error decreases as we train the model for more iterations., and plateaus after 3 iterations. Model generalizes well to new shapes and new poses. (B) Alignment error declines as we use more PCs. The improvement after 12th PC is marginal. (C) Fetal model outperforms baseline models.

then registered to the clinical cohort to measure anthropometric measurements. To quantitatively evaluate the alignment, we compute the one-sided median Chamfer distance from the segmentation surface to the model. To derive a metric for time series, we average the median Chamfer distances over time.

Baselines. While our closest baseline model is infant SMPL [7], it may not effectively capture fetal body shape since newborns are typically larger than fetuses and their bodies have different proportions. We downscale the newborn model to 75% of its size to match the fetal population [7,17]. We evaluate two variants of the infant SMPL model. The first variant — infant SMPL — jointly estimates β and θ without an unpose step. The second variant — infant SMPL w. unpose — decouples the estimation of β and θ with an unposing step, following our registration method.

Quantitative Results. Fig. 2A reports the alignment error as a function of training iterations. The model achieved a median alignment error of 3.01 mm on the training dataset in the third iteration, after which the performance plateaued. Similarly, in the “New Pose” and “New Shape” datasets, the model achieves errors of 3.14 mm and 3.17 mm. These results indicate that our model generalizes well to unseen poses and shapes. We use the 3rd iteration model in the following experiments. Fig. 2B reports the alignment error as a function of number of principal components (PCs). The alignment error begins to converge at approximately 8 PCs. Interestingly, while including more PCs enables finer shape adjustments, it results in a qualitatively less smooth body surface. To balance the quality of fit and that of visualization, we use the first 10 PCs, which capture 98.3% of the variance, consistently with SMPL and SMIL.

Fig. 2C compares our method with the two baseline methods on the “New Shape” dataset. Across all subjects, infant SMPL model achieves a median alignment error of 5.21 mm. We observe that if the shape parameters β and pose parameters θ are estimated jointly, the predicted fetal body shapes are typically smaller (infant SMPL). This is likely because in the fetal pose, the body parts are clumped together. A smaller predicted body shape results in lower alignment




Fig. 3. Inter-subject shape differences and population-level shape PCs. (A) The first 3 PCs describe body size, body thickness, and ratio between leg length and body length. PC1 varies from -2 to $+2$ standard deviations from left to right, while PC2 and PC3 vary from -4 to $+4$ to make shape differences more visible. (B) Scatter plot of the first two PCs with S0, S1, and S2 subjects marked in the distribution. (C) The MR images and posed body shapes of S0, S1, and S2. (D) Side by side visualization of S2+, S2, S0, S1, and S1+.

error but does not reflect the true body shape. If we decouple the shape and pose estimation using the “unpose” step, the body alignment improves significantly to 3.61 mm (infant SMPL w. unpose). Finally, our model with fully decoupled estimation achieved the lowest alignment error of 3.07 mm.

We train the final model using the entire research cohort to fully utilize our data. We then align the model to all research subjects to evaluate clinically relevant anthropometric measurements. The model-derived metrics are strongly correlated with gestational age (GA) (HC: $R = 0.79$, AC: $R = 0.80$, FL: $R = 0.77$, BL: $R = 0.78$, BV: $R = 0.82$, $p < 0.001$ for all measurements) consistent with prior literature [17]. We align the model to the clinical cohort, and the median segmentation-to-model Chamfer distance is 4.2 mm.

Visualization. In Fig. 3A, we visualize population-level shape characteristics using the learned eigen-shapes. The first three PCs respectively capture body size, abdominal thickness, and ratio of leg length to body length. In Fig. 3C, we visualize and compare canonical shapes from three subjects (S0, S1, and S2). We choose S0 to be the one closest to the mean shape. We select S1 and S2 based on their deviations in the second principal component. We ensure

they have small deviations in the first PC to isolate the effects of the second PC. Shape differences among S1, S2, and S0 are visually subtle (Fig. 3D). To amplify these differences, we magnify the deviations of S1 and S2 relative to S0 by a factor of two (Fig. 3D). The resulting magnified *S1+* and *S2+* shapes show clear differences in body thickness compared to S0. Our tool enables the visualization and amplification of subtle shape differences among subjects for improved comparison.

4 Conclusions

In this paper, we presented a 3D articulated statistical fetal body model. We trained the model from noisy segmentations and keypoints extracted from the MRI time series. A key enabler is alternating between the estimation of the shape and pose parameter. The resulting model accurately explains unseen fetal shapes. Moreover, the model provides clinically relevant measurements that are traditionally difficult to obtain from segmentations and keypoints. This work promises improved fetal development assessment and prenatal care.

Acknowledgements. This work has been funded by NIH NIBIB 1R01EB032708, NIH NICHD 1R01HD114338, NIH NIBIB 1R01EB036945, and MIT CSAIL-Wistron Program. We are grateful to Neerav Karani, Neel Dey, Marilyn Keller, Nikolas Hesse, Sergi Pujades, Maz Abulnaga, and Justin Solomon for their insightful discussions.

Disclosure of Interests. The authors have no competing interests to declare.

References

1. Alldieck, T., Magnor, M., Xu, W., Theobalt, C., Pons-Moll, G.: Video based reconstruction of 3d people models. In: IEEE Conference on Computer Vision and Pattern Recognition (2018)
2. Barron, J.T.: A general and adaptive robust loss function. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4331–4339 (2019)
3. Billot, B., Dey, N., Moyer, D., Hoffmann, M., Turk, E.A., Gagoski, B., Grant, P.E., Golland, P.: SE(3)-equivariant and noise-invariant 3D rigid motion tracking in brain MRI. IEEE Transactions on Medical Imaging (2024)
4. Billot, B., Dey, N., Turk, E.A., Grant, E., Golland, P.: Network conditioning for synergistic learning on partial annotations. In: Medical Imaging with Deep Learning (2024)
5. Candès, E.J., Li, X., Ma, Y., Wright, J.: Robust principal component analysis? Journal of the ACM (JACM) **58**(3), 1–37 (2011)
6. Chi, Z., Cong, Z., Wang, C.J., Liu, Y., Turk, E.A., Grant, P.E., Abulnaga, S.M., Golland, P., Dey, N.: Dynamic neural fields for learning atlases of 4d fetal mri time-series. arXiv preprint arXiv:2311.02874 (2023)

7. Hesse, N., Pujades, S., Black, M.J., Arens, M., Hofmann, U.G., Schroeder, A.S.: Learning and tracking the 3d body shape of freely moving infants from rgb-d sequences. *IEEE transactions on pattern analysis and machine intelligence* **42**(10), 2540–2551 (2019)
8. Hesse, N., Pujades, S., Romero, J., Black, M.J., Bodensteiner, C., Arens, M., Hofmann, U.G., Tacke, U., Hadders-Algra, M., Weinberger, R., et al.: Learning an infant body model from rgb-d data for accurate full body motion analysis. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st International Conference, Granada, Spain, September 16–20, 2018, Proceedings, Part I*. pp. 792–800. Springer (2018)
9. Kaufmann, M., Vechev, V., Mylonopoulos, D.: aitviewer (7 2022). <https://doi.org/10.5281/zenodo.10013305>, <https://github.com/eth-ait/aitviewer>
10. Kingma, D.P.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
11. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: Smpl: A skinned multi-person linear model. In: *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pp. 851–866 (2023)
12. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3d surface construction algorithm. In: *Seminal graphics: pioneering efforts that shaped the field*, pp. 347–353 (1998)
13. Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3d hands, face, and body from a single image. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10975–10985 (2019)
14. Ravi, N., Reizenstein, J., Novotny, D., Gordon, T., Lo, W.Y., Johnson, J., Gkioxari, G.: Accelerating 3d deep learning with pytorch3d. *arXiv:2007.08501* (2020)
15. Sadhwani, A., Wypij, D., Rofeberg, V., Gholipour, A., Mittleman, M., Rohde, J., Velasco-Annis, C., Calderon, J., Friedman, K.G., Tworetzky, W., et al.: Fetal brain volume predicts neurodevelopment in congenital heart disease. *Circulation* **145**(15), 1108–1119 (2022)
16. Sajapala, S., AboEllail, M.A.M., Kanenishi, K., Mori, N., Marumo, G., Hata, T.: 4d ultrasound study of fetal movement early in the second trimester of pregnancy. *Journal of Perinatal Medicine* **45**(6), 737–743 (2017)
17. Snijders, R., Nicolaides, K.: Fetal biometry at 14–40 weeks’ gestation. *Ultrasound in Obstetrics and Gynecology: The Official Journal of the International Society of Ultrasound in Obstetrics and Gynecology* **4**(1), 34–48 (1994)
18. Uus, A.U., Hall, M., Grigorescu, I., Avena Zampieri, C., Egloff Collado, A., Payette, K., Matthew, J., Kyriakopoulou, V., Hajnal, J.V., Hutter, J., et al.: Automated body organ segmentation, volumetry and population-averaged atlas for 3d motion-corrected t2-weighted fetal body mri. *Scientific Reports* **14**(1), 6637 (2024)
19. Vasung, L., Xu, J., Abaci-Turk, E., Zhou, C., Holland, E., Barth, W.H., Barnewolt, C., Connolly, S., Estroff, J., Golland, P., et al.: Cross-sectional observational study of typical in utero fetal movements using machine learning. *Developmental neuroscience* **45**(3), 105–114 (2023)
20. Xu, J., Turk, E.A., Gagoski, B., Golland, P., Grant, P.E., Adalsteinsson, E.: Motion analysis in fetal mri using deep pose estimator. In: *Proceedings of the International Society for Magnetic Resonance in Medicine... Scientific Meeting and Exhibition. International Society for Magnetic Resonance in Medicine. Scientific Meeting and Exhibition. vol. 29. NIH Public Access* (2021)

21. Xu, J., Zhang, M., Turk, E.A., Grant, P.E., Golland, P., Adalsteinsson, E.: 3d fetal pose estimation with adaptive variance and conditional generative adversarial network. In: Medical Ultrasound, and Preterm, Perinatal and Paediatric Image Analysis: First International Workshop, ASMUS 2020, and 5th International Workshop, PIPPI 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, October 4-8, 2020, Proceedings 1. pp. 201–210. Springer (2020)
22. Xu, J., Zhang, M., Turk, E.A., Zhang, L., Grant, P.E., Ying, K., Golland, P., Adalsteinsson, E.: Fetal pose estimation in volumetric mri using a 3d convolution neural network. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part IV 22. pp. 403–410. Springer (2019)