Automatic fetal fat quantification from MRI

Netanell Avisdris* ✉[1,2][0000−0002−0719−1327], Aviad Rabinowich*[2,3,4], Daniel Fridkin[1], Ayala Zilberman[4,5], Sapir Lazar[3,4], Jacky Herzlich[4,6], Zeev Hananis[2], Daphna Link-Sourani[4], Liat Ben-Sira[3,4], Liran Hiersch[4,5], Dafna Ben Bashat[2,4,7], and Leo Joskowicz[1]

[1] School of Computer Science and Engineering, The Hebrew University of Jerusalem, Israel {netana03,josko}@cs.huji.ac.il
[2] Sagol Brain Institute, Tel Aviv Medical Center, Israel
[3] Department of Radiology, Tel Aviv Medical Center, Tel Aviv, Israel
[4] Sackler Faculty of Medicine, Tel Aviv University, Tel Aviv, Israel
[5] Department of Obstetrics and Gynecology, Lis Hospital for Women, Tel Aviv Sourasky Medical Center, Tel Aviv, Israel
[6] Neonatal Intensive Care Unit, Dana Dwek Children's Hospital, Tel Aviv Sourasky Medical Center, Tel Aviv, Israel
[7] Sagol School of Neuroscience, Tel Aviv University, Tel Aviv, Israel

Abstract. Normal fetal adipose tissue (AT) development is essential for perinatal well-being. AT, or simply fat, stores energy in the form of lipids. Malnourishment may result in excessive or depleted adiposity. Although previous studies showed a correlation between the amount of AT and perinatal outcome, prenatal assessment of AT is limited by lacking quantitative methods. Using magnetic resonance imaging (MRI), 3D fat- and water-only images of the entire fetus can be obtained from two-point Dixon images to enable AT lipid quantification. This paper is the first to present a methodology for developing a deep learning (DL) based method for fetal fat segmentation based on Dixon MRI. It optimizes radiologists' manual fetal fat delineation time to produce annotated training dataset. It consists of two steps: 1) model-based semi-automatic fetal fat segmentations, reviewed and corrected by a radiologist; 2) automatic fetal fat segmentation using DL networks trained on the resulting annotated dataset. Segmentation of 51 fetuses was performed with the semi-automatic method. Three DL networks were trained. We show a significant improvement in segmentation times (3:38 hours → < 1 hour) and observer variability (Dice of 0.738 → 0.906) compared to manual segmentation. Automatic segmentation of 24 test cases with the 3D Residual U-Net, nn-UNet and SWIN-UNetR transformer networks yields a mean Dice score of 0.863, 0.787 and 0.856, respectively. These results are better than the manual observer variability, and comparable to automatic adult and pediatric fat segmentation. A radiologist reviewed and corrected six new independent cases segmented using the best performing network (3D Residual U-Net), resulting in a Dice score of 0.961 and a significantly reduced correction time of 15:20 minutes. Using these novel segmentation methods and short MRI acquisition time, whole body subcutaneous lipids can be quantified for individual fetuses in the clinic and large-cohort research.

1 Introduction

The adipose tissue (AT) is essential for fetal development and reflects the overall fetal energy balance. AT, or simply fat, stores energy in the form of lipids. Typically, well-nourished fetuses show an accelerated AT growth from the 28 weeks of gestation onward [20]. Although birth weight is often used as a proxy of fetal nutrition and as a predictor of adverse perinatal outcome, prior studies suggest that fetal body fat may show high correlation and be better predictor for neonatal outcomes [1,4,27]. Previous studies using ultrasound (US) showed alternations of fetal AT related to fetal growth restriction and excess of fetal AT in cases of maternal diabetes [11,18,25]. Furthermore, these changes also correlate with neonatal outcomes, emphasizing the clinical relevance of AT quantification. Although US is the method of choice for fetal development assessment, it is limited by the lack of true 3D information. Thus, fat estimation currently relies on linear and area measurements of selected fetal body parts, e.g., the abdomen and the limbs, and on estimated fractional limb volume [19], with no full body AT fat content volume quantification and analysis.

Magnetic resonance imaging (MRI) provides 3D multi-contrast information that indirectly characterizes the microstructural properties of tissue. For fat assessment, the two-point Dixon method [8], a proton chemical shift MRI technique that produces separated fat-only and water-only images from a dual-echo acquisition, is used. In this method, water and fat signals alternate between being summed and subtracted, yielding fat-only and water-only images that can be analyzed to quantify lipids. This method is used to quantify the lipid content of various organs, most extensively for hepatosteatosis assessment [13].

Fetal MRI is increasingly used as a complementary method to US, mainly for detecting central nervous system (CNS) and non-CNS anomalies, including thoracic, gastrointestinal, genitourinary, and skeletal anomalies [5]. A few studies measured AT of fetuses with normal-growth and with maternal diabetes [2,3,12]. However, these studies are limited, as they rely on linear measurements [2], on local assessment [3], or require laborious manual segmentation [12].

Automatic 3D MRI segmentation can enable accurate and reliable routine AT lipid quantification. Recently, deep learning (DL) based models have been increasingly used for automated segmentation of structures in medical imaging, including fetal structures [23,29]. While most methods address fetal brain segmentation, a few have been developed for other fetal structures and organs.

To the best of our knowledge, no automatic or semi-automatic methods for fetal fat quantification in US or MRI have been developed. Roelants et al. [24] described automated methods for US limb soft-tissue quantification, a surrogate for subcutaneous AT lipid deposition. Mack et al. [22] presented a semi-automatic method for fractional limb volume assessment and showed that it reduces observer variability and annotation time. Recent papers described automatic DL methods for fat segmentation in Dixon scans of adults in the pancreas [21] and

visceral and subcutaneous fat [10,17]. Estrada et al. [10] reported that manual fat delineation is tedious and time-consuming, thus limiting its clinical and research applications. Fat segmentation poses several significant challenges compared to other structures. First, ground-truth manual fetal fat delineations are time-consuming, laborious and difficult to acquire [10,12], and suffer from high observer variability [15]. Indeed, fat structure is thin and complex, is distributed in various locations, and may appear as multiple disjoint components with a wide variety of shapes. Moreover, Dixon scans exhibit inherent limitations, e.g., signal inhomogeneity, artifacts due to high sensitivity to fetal and maternal motion, and obscured lipid-poor tissues in fat-only images, which lack the fetus context without the water-only images and requires additional MRI sequences.

In this paper, we present a methodology for creating the first reported automatic DL-based method for fetal fat segmentation on Dixon images. Its contributions are three-fold: (1) a semi-automatic segmentation method for fetal fat delineation that substantially shortens the manual annotation time and reduces inter-observer variability; (2) training and evaluation of three state-of-the-art deep learning models for fetal fat segmentation on the validated ground-truth segmentations generated by the semi-automatic method; (3) quantification of the manual and semi-automatic delineation observer variability and annotation time of fetal fat.

2   Methodology

Our methodology for fetal fat segmentation consists of: 1) semi-automatic fetal fat model-based segmentation whose aim is to shorten the manual fetal fat delineation time required to produce annotated training data, and; 2) automatic fetal fat segmentation with DL networks trained on the resulting annotated dataset.

2.1   Semi-automatic fetal AT segmentation

The inputs for the semi-automatic fetal fat segmentation method are fetal body TRUFI and fat-only Dixon MRI sequences. The output is an initial subcutaneous fetal fat segmentation on the Dixon scan (Fig. 1). The TRUFI sequence is used to obtain fetal body segmentation that defines the Volume of Interest (VOI) in the Dixon scan on which the fetal fat segmentation is computed. Although the TRUFI and Dixon scans are acquired at subsequent times, they have different field-of-views and may not be aligned due to fetal and maternal motion.

First, the fetal body is automatically segmented on the TRUFI scan with the DL model described in [9]. The resulting segmentation mask is then mapped to the Dixon scan using the scanning position information to define a VOI that includes the entire fetus and excludes the maternal abdominal fat regions. Manual adjustments to the resulting VOI are performed as required. The fetal fat within the VOI is then thresholded with a pre-defined value chosen experimentally. Connected components with < 50 voxels are discarded to remove artifacts with similar fat intensity. The result is then reviewed and corrected by an expert radiologist to produce high-quality, validated ground truth segmentation masks.
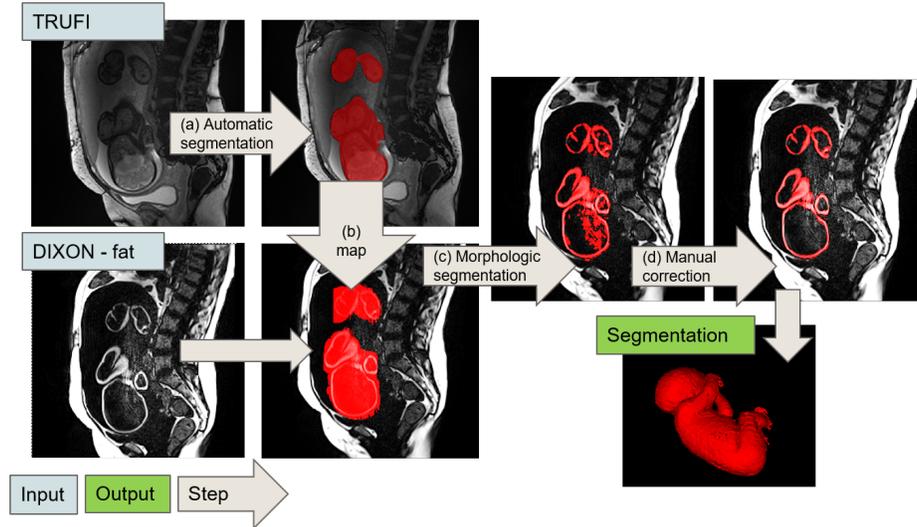
Fig. 1. Semi-automatic fetal fat segmentation method. The inputs (light blue) are fetal body TRUFI and fat-only Dixon scans. The four steps (grey arrows) of the method are: (a) automatic segmentation of the fetal body in the TRUFI scans; (b) mapping of the resulting segmentation as a VOI on the fat-only Dixon volume; (c) automatic segmentation of the fetal fat on the Dixon scan by thresholding and morphology operations; (d) revision and manual correction by a radiologist of the resulting fetal fat segmentation mask on each slice. The output (green) is a validated fetal fat segmentation mask.

2.2 Automatic fetal fat segmentation

Automatic fetal fat segmentation is performed with a state-of-the-art DL model trained on the high-quality annotations of the fetal fat created in the previous step. The inputs are fat-only and water-only Dixon volumes, which are acquired simultaneously, and are thus aligned. The output is the fetal fat segmentation.

We evaluate three models: (1) Residual 3D U-Net [16]; (2) nn-UNet [14], and; (3) SWIN-UNetR [28]. The 3D U-Net and nn-UNet are fully convolutional (FCN) encoder-decoder networks, where the decoder network is connected to the encoder network through skip connections. SWIN-UNetR is a combined transformer-FCN network. We briefly describe them next.

Residual 3D U-Net is a 3D-UNet [6] in which the encoder and decoder network blocks are residual units. The 3D U-Net is a FCN which consists of a contraction and an expansion path. The contraction path is an encoder that captures the context, while the expansion path is a decoder network that performs up-sampling to recover the segmentation map size. The encoder and decoder paths are connected by path skip connections to share localization information.

nn-UNet is a biomedical image segmentation framework that automatically adapts to a dataset characteristics. It selects a segmentation network from a wide range of options, including 2D- and 3D- U-Net, and configures their hyper-
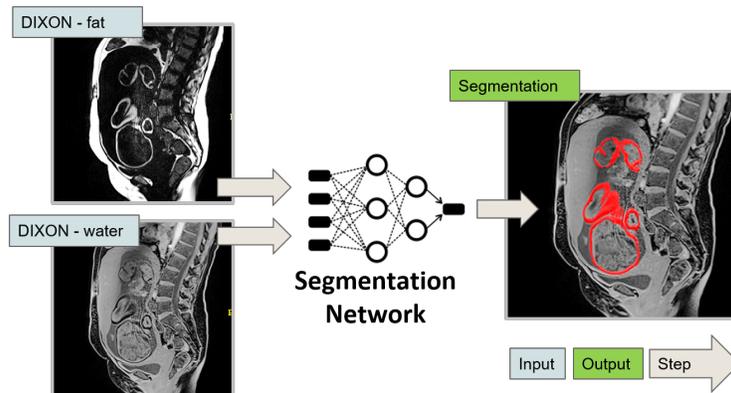
Fig. 2. Automatic fetal fat segmentation. The inputs (light blue) are fat-only and water-only Dixon scans. The output (green) is the fetal fat segmentation with a DL network.

parameters, e.g., patch size, batch size, or learning rate. It achieves state-of-the-art results on a wide range of biomedical image segmentation scenarios.

SWIN-UNetR: Transformers are new DL architectures that achieved state-of-the-art results in a wide range of machine learning tasks, including medical imaging [26]. Shifted windows (SWIN) Transformers is a hierarchical vision transformer that allows for local computing of self-attention with non-overlapping windows. SWIN is more efficient than regular vision transformers and is well-suited for tasks requiring multi-scale modeling due to its hierarchical nature.is a vision transformer for biomedical image analysis. It consists of a SWIN transformer as the encoder and a CNN-based decoder. SWIN-UNetR computes self-attention in an efficient shifted window partitioning scheme. It is currently the best performing model for a wide variety of biomedical image segmentation tasks.

Implementation details: The 3D-UNet and SWIN-UNetR models were implemented in MONAI [7]; nn-UNet was implemented in Pytorch. All networks were trained on 128x128x128 patches for 300 epochs on a single NVIDIA V100 GPU. The 3D-UNet was trained with an ADAM optimizer with an initial learning rate of 1×10^{-2} and batch size of 10. SWIN-UNetR was trained with an AdamW optimizer with a warm-up cosine scheduler of 50 iterations and batch size of 1 patch. Both 3D-UNet and SWIN-UNetR are trained with random patch cropping and Gaussian noise augmentations, with Dice loss. Since nn-UNet is an automated framework, no training hyper-parameters were selected. For inference, we use a sliding window with an overlap of 0.7 for neighboring voxels.

3   Experimental Results

We conducted three studies. Study 1 quantifies the observer variability of manual and semi-automatic fat delineation. The manual delineation observer variabil-

	Manual				Semi-automatic			
	mean	std	min	max	mean	std	min	max
Dice	0.738	0.092	0.567	0.865	**0.906**	**0.084**	0.744	0.981
Hausdorff [mm]	21.09	24.66	2.24	86.82	**16.88**	**5.10**	10.44	28.32
ASSD [mm]	1.90	4.11	0.23	13.57	**0.36**	**0.28**	0.07	0.87
VD [mL]	-	-	-	-	18.69	20.63	2.50	73.48
RVD [%]	29.91	12.54	6.26	46.34	**9.26**	**9.14**	1.10	26.50

Table 1. Inter-observer variability of fetal fat delineation between two radiologists using two methods: manual and semi-automatic. Bold face indicates best results for each of the five metrics.

ity provides a reference for the expected target accuracy; the semi-automatic delineation observer variability quantifies the expected improvement in observer agreement. The study also measures the radiologist time required for each method. Study 2 quantifies the performance of the three automatic DL fetal fat segmentation methods. Study 3 quantifies and analyze the segmentation accuracy and the correction time using the best-performing automatic DL network.

Study population: Retrospective fetal MRI studies were collected at the Tel Aviv Sourasky Medical Center, Israel, between 2019 and 2022. The Institutional Review Board approved the study and waived the need for informed consent. The dataset consisted of 57 singleton fetuses ranging between 31 and 39+1 gestation weeks (mean = 33.9, std = 1.8). Participants were referred to fetal MRI for various clinical indications. Cases with chromosomal or congenital anomalies were excluded.

Data acquisition and ground truth generation: Patients were scanned on one of three 3T MRI scanners (Skyra, Prisma and Vida, Siemens Healthineers). Two subsequent sequences were used for this study: (a) free-breathing T2-weighted TRUFI sequence with voxel resolution of $0.78 \times 0.78 \times 2mm^3$ and an acquisition time of 60 seconds, and (b) two-point Dixon sequence with voxel resolution of $1.25-1.4 \times 1.25-1.4 \times 1.5-2mm^3$. The Dixon sequence was acquired with a single breath-hold, with acquisition time of 18-20 seconds. ITK-SNAP(v. 3.8.0) [30] was used for manual delineation and segmentation correction.

Evaluation metrics: Five metrics were used to estimate observer agreement and evaluate segmentation methods: Dice similarity coefficient, Hausdorff distance, Average Symmetric Surface Distance (ASSD), volume difference (VD), and relative volume difference (RVD). RVD is defined as the VD divided by the fetal body volume. We used a two-sided t-test to estimate statistical significance; $p < 0.05$ was considered significant.

3.1   Study 1: Manual and semi-automatic observer variability

No previous study has analyzed the inter-observer variability in fetal fat segmentation. Manual and semi-automatic segmentations were assessed by two radiol-

	3D Residual UNet				nn-UNet				SWIN-UNetR			
	mean	std	min	max	mean	std	min	max	mean	std	min	max
Dice	0.862	0.063	0.609	0.926	0.787	0.169	0.269	0.927	0.856	0.082	0.538	0.933
Hausdorff [mm]	18.05	4.88	9.70	28.58	37.19	34.06	9.43	134.85	20.63	5.74	9.17	36.07
ASSD [mm]	0.60	0.33	0.30	1.93	1.54	2.09	0.32	7.52	0.66	0.51	0.25	2.78
VD [mL]	16.93	27.16	0.74	138.38	27.67	36.08	1.53	177.89	19.45	38.47	0.19	189.00
RVD [%]	7.11	7.44	0.50	36.66	13.59	15.78	0.82	71.60	8.36	11.50	0.12	50.07

Table 2. Study 2 results: comparison of three networks for fetal AT segmentation: 3D Residual UNet, nn-UNet and SWIN-UNetR on test-set (n=24). Bold face indicates best results for a metric.

ogists that were blinded to the study indication and to the gestational age. A previous study noted extremely long manual segmentation times [12]. Therefore, four consecutive slices on 10 scans were randomly selected for manual segmentation in the fetal volume. For the semi-automatic segmentation, the first observer (R1) segmented the entire cohort. Ten cases were assigned to a second observer (R2) to establish a baseline for the automatic method and to assess the segmentation quality. Manual delineation and correction times were recorded.

Table 1 lists the results. Overall, manual fat segmentation was very time-consuming, with a mean of 3:38 hours (R1: 3:43 hours, R2: 3:34 hours) for the entire fat fetal volume, similar to the reported time in [12]. Semi-automatic segmentation significantly reduced total segmentation time to a mean of 0:57 hours ($p = 6.3 \times 10^{-7}$). Furthermore, semi-automatic delineation significantly increased the inter-observer agreement for Dice and RVD: Dice increased by 0.186 ($p = 3 \times 10^{-5}$). Also, Hausdorff decreased by $4.20mm$ ($p = 0.585$), ASSD decreased by $1.54mm$ ($p = 0.228$) and RVD decreased by 20.55% ($p = 3 \times 10^{-5}$).

3.2   Study 2: Automatic fetal AT segmentation

We evaluate three networks for the automatic fetal AT segmentation task: 3D Residual U-Net, nn-UNet, and SWIN-UNetR. A dataset of 51 manually corrected volumes from semi-automatic fetal fat segmentation was used. The networks were trained on 21 volumes; six volumes were used for validation and hyper-parameter choice. The networks were tested on an independent set of 24 volumes. Results are listed on Table 2. Overall, the 3D Residual UNet yielded the best results, achieving a mean Dice score of 0.862, above the manual delineation observer variability. Fig. 3 shows illustrative fetal fat segmentation results of three fetuses.

3.3   Study 3: Analysis of manual corrections after automatic segmentation

To evaluate the effectiveness of the automatic method, six additional test cases were segmented using the 3D Residual UNet, which was the best performing
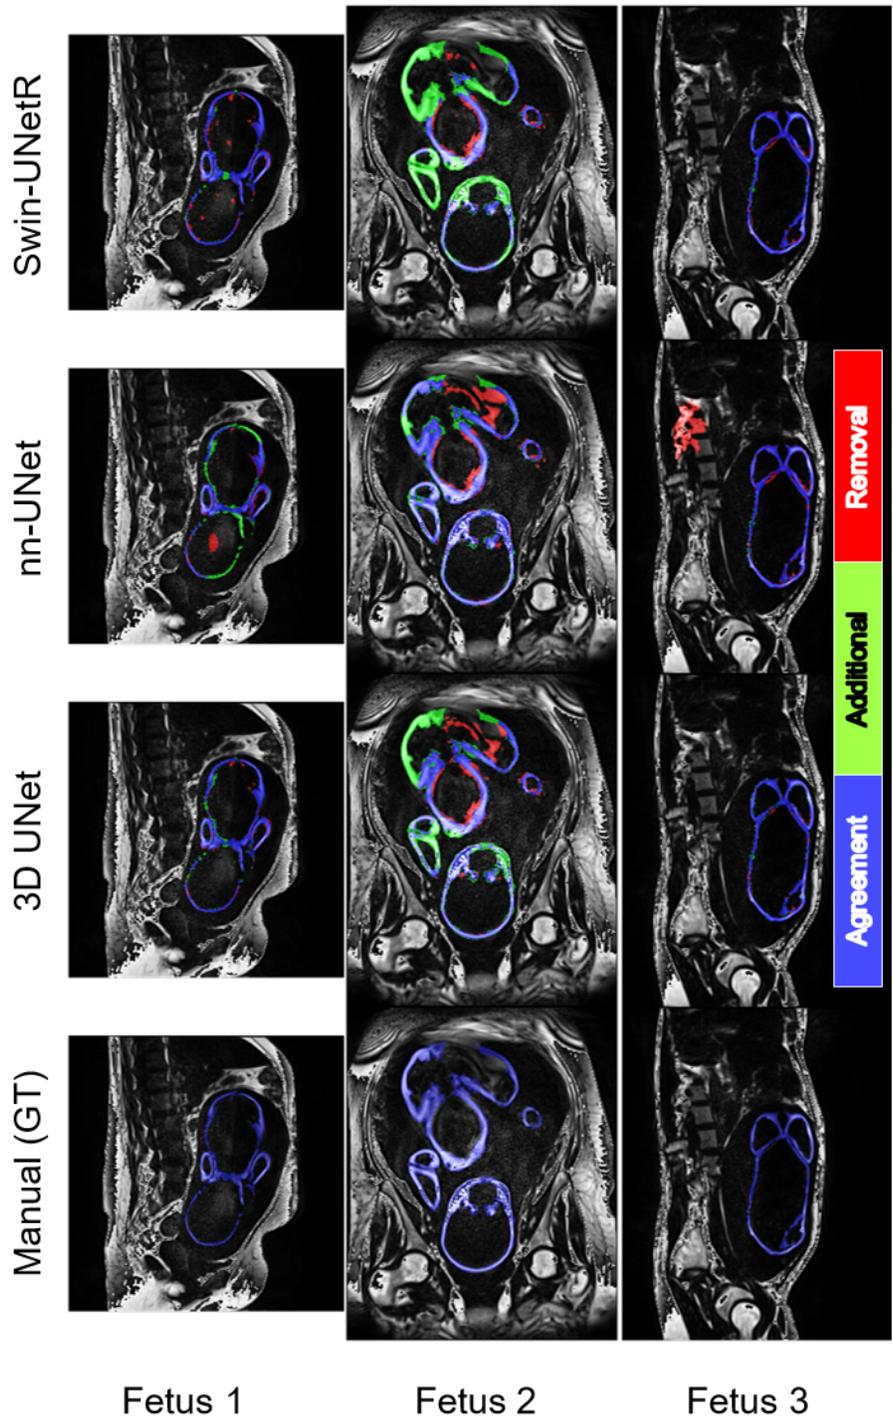
Fig. 3. Comparison of automatic segmentation networks performance. Three examples (columns) of manual ground-truth and three automatic segmentation networks (rows). The manual and automatic segmentation agreement voxels are shown in blue. Undersegmentation is shown in green. Oversegmentation is shown in red. The nn-UNet tends to yield over-segmented components, including large out-of-fetus components (Fetus 3).

	mean	std	min	max
Dice	0.961	0.025	0.915	0.983
Hausdorff [mm]	48.72	29.67	13.08	84.30
ASSD [mm]	0.59	1.00	0.10	2.63
VD [mL]	7.54	10.91	1.62	29.35
RVD [%]	4.68	5.48	1.37	15.56

Table 3. Study 3 results. Manual correction on best performing automatic segmentation network (3D Residual UNet) for 6 cases.

network. The resulting segmentations were manually corrected by a radiologist (R1). The manual correction times were recorded.

Results are listed in Table 3. Overall, only minor revisions were required (Dice of 0.961), and correction times were reduced to an average of 15:20 minutes. However, the Hausdorff distance was relatively high, 48.72mm, compared to the inter-observer one (16.88mm). Visual examples of corrections are presented in Figure 4. Note that minor corrections were performed on small, out-of-fetus voxels (yellow arrow on figure 4), which may explain this result. Additional corrections were performed on lipid-poor areas such as the scalp, intensity artifacts, and non-subcutaneous fat depots (e.g. , perirenal fat). All fetus body parts were segmented using the automatic method, with only minor voxels difference, resulting in a relatively small RVD of 4.68%.

4   Discussion

In this paper, we describe methods for semi-automatic and automatic fetal subcutaneous fat segmentation. We propose a semi-automatic segmentation method to delineate the fetal subcutaneous fat and showed that it significantly improves segmentation times and observer variability compared to manual segmentation. We then used the resulting segmentation masks to train an automatic deep-learning network, which yields accurate results, better than the inter-observer variability on manual segmentation, and shorten the correction time.

Automatic fat segmentation, and specifically fetal fat segmentation, is known to be a very challenging task. Previous works suggest that segmentation observer variability can be used as a surrogate to estimate the expected accuracy of automatic segmentation [15]. Here we show that fetal subcutaneous fat delineation has a high inter-observer variability with a Dice of 0.738 vs. 0.85-0.95 typically observed for other fetal structures [29], which may be correlated to the difficulty of automatic segmentation. Estrada et al. [10] report an inter-observer variability of 0.982 and 0.788 Dice scores for adult subcutaneous and visceral fat, respectively. Adult visceral fat segmentation problem is similar to that of fetal subcutaneous fat, as it is a complex, sparse structure with many connected components, due to the intestines and the large surface it covers.
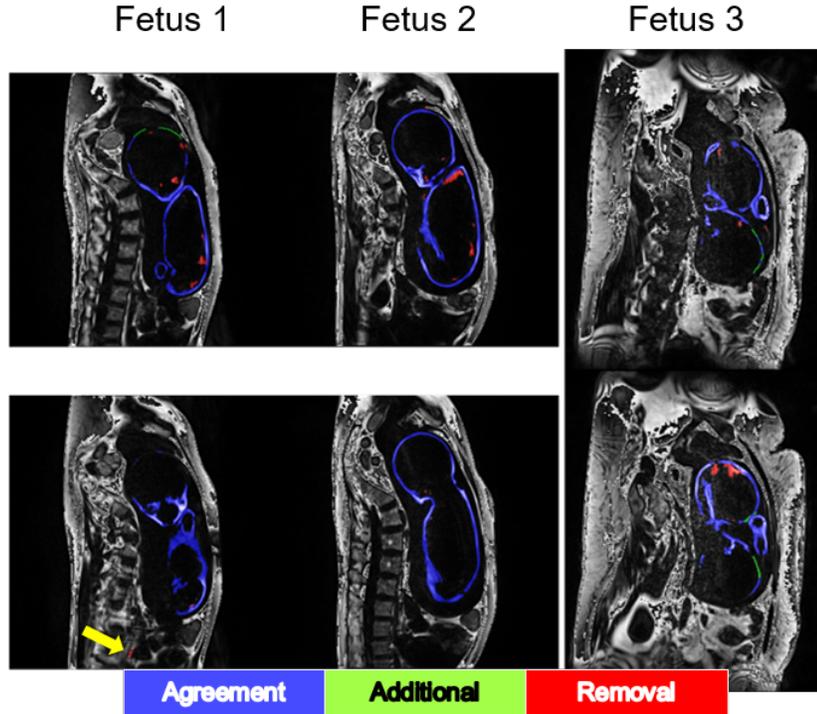
Fig. 4. Examples of manual correction on best performing automatic segmentation network (3D Residual UNet) of three fetuses (columns) from two representative slices (rows). The yellow arrow points to a small out-of-fetus segmented compartment that needs to be erased.

To address these issues, we used a semi-automatic method that reduced the observer variability and shorten segmentation times. These results suggest the advantage of using a semi-automatic segmentation to more efficiently produce more accurate ground-truth annotations used to train an automatic DL method. This scheme is similar to the one described in Kway et al. [17] for pediatric subcutaneous and visceral fat segmentations.

In this work, we explore the use of three neural networks for fetal fat segmentation. In particular, we compare the state-of-the-art models Residual 3D UNet, nn-UNet [14] and SWIN-UNetR [28]. In previous works [10,17], variants of 3D-UNet were used to automate fat segmentation. However, nn-UNet and SWIN-UNetR, which yield state-of-the-art results in a variety of biomedical image segmentation scenarios were not evaluated on fat segmentation tasks. In our study, SWIN-UNetR and 3D-UNet achieved similar results, while nn-UNet performed poorly. A possible explanation might relate to the different characteristics of fat segmentation, which differ from what nn-UNet was designed for [14]. Related fat segmentation methods achieved a Dice of 0.850 and 0.872 for

pediatric and adults, respectively [10,17]. Our result of 0.862 is comparable to those results, and is better than observer variability of manual delineation. Moreover, the automatic method achieved an RVD of 7.11%, which is lower than the interobserver variability RVD of 29.91% (manual) and 9.26% (semi-automatic).

Our study have several limitations. First, this is a single center cohort, which may result in poor generalization. Second, we include fetuses with a gestation age of 31 weeks or more, which may limit the automatic segmentation for younger fetuses. However, previous studies showed that the fetal fat is apparent in MRI only from the 28^{th} week with rapid third trimester lipid accumulation [3,12]. Therefore, the pivotal age for fat quantification is similar to that of our study. Third, our cohort size of 57 fetuses is relatively small. Future studies should use a larger, and wider gestational age range to develop and assess the applicability of automatic methods for fetal fat quantification. Lastly, we explored the naive use of state-of-the-art DL methods. Future studies may design a more targeted solution of DL method in order to further improve fetal fat automatic segmentation and reduce the need for manual correction.

5   Conclusion

Here we propose the first method to automate subcutaneous fetal fat segmentation. We show that using state-of-the-art segmentation methods and short acquisition time MRI sequences, whole fetal body subcutaneous lipids can be quantified. We anticipate that our method can be used to study normal changes of fetal fat with gestational age and its relation to various abnormal conditions, e.g., gestational diabetes and fetal growth restriction.

References

1. Banting, S.A., Dane, K.M., Charlton, J.K., Tong, S., Hui, L., Middleton, A.L., Gibson, L.K., Walker, S.P., MacDonald, T.M.: Estimation of neonatal body fat percentage predicts neonatal hypothermia better than birthweight centile. The Journal of Maternal-Fetal & Neonatal Medicine pp. 1–8 (2022)
2. Berger-Kulemann, V., Brugger, P.C., Reisegger, M., Klein, K., Hachemian, N., Koelblinger, C., Weber, M., Prayer, D.: Quantification of the subcutaneous fat layer with mri in fetuses of healthy mothers with no underlying metabolic disease vs. fetuses of diabetic and obese mothers. J. Perinat. Med. (2012)
3. Blondiaux, E., Chougar, L., Gelot, A., Valence, S., Audureau, E., Ducou le Pointe, H., Jouannic, J.M., Dhombres, F., Garel, C.: Developmental patterns of fetal fat and corresponding signal on t1-weighted magnetic resonance imaging. Pediatr. Radiol. **48**(3), 317–324 (2018)
4. Carberry, A.E., Raynes-Greenow, C.H., Turner, R.M., Askie, L.M., Jeffery, H.E.: Is body fat percentage a better measure of undernutrition in newborns than birth weight percentiles? Pediatric research **74**(6), 730–736 (2013)

5. Cassart, M., Garel, C.: European overview of current practice of fetal imaging by pediatric radiologists: a new task force is launched. Pediatr. Radiol. **50**(12), 1794–1798 (2020)
6. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3d u-net: learning dense volumetric segmentation from sparse annotation. In: Proc Int Conf on Medical Image Computing and Computer Assisted Intervention. pp. 424–432. Springer (2016)
7. Consortium, M.: Monai: Medical open network for ai (Jun 2022). https://doi.org/10.5281/zenodo.6639453
8. Dixon, W.T.: Simple proton spectroscopic imaging. Radiology **153**(1), 189–194 (1984)
9. Dudovitch, G., Link-Sourani, D., Sira, L.B., Miller, E., Bashat, D.B., Joskowicz, L.: Deep learning automatic fetal structures segmentation in mri scans with few annotated datasets. In: Proc Int Conf on Medical Image Computing and Computer Assisted Intervention (2020)
10. Estrada, S., Lu, R., Conjeti, S., Orozco-Ruiz, X., Panos-Willuhn, J., Breteler, M.M., Reuter, M.: Fatsegnet: A fully automated deep learning pipeline for adipose tissue segmentation on abdominal dixon mri. Magn Reson Med **83**(4), 1471–1483 (2020)
11. Gardeil, F., Greene, R., Stuart, B., Turner, M.J.: Subcutaneous fat in the fetal abdomen as a predictor of growth restriction. Obstetrics & Gynecology **94**(2), 209–212 (1999)
12. Giza, S.A., Koreman, T.L., Sethi, S., Miller, M.R., Penava, D.A., Eastabrook, G.D., McKenzie, C.A., de Vrijer, B.: Water-fat magnetic resonance imaging of adipose tissue compartments in the normal third trimester fetus. Pediatr. Radiol. **51**(7), 1214–1222 (2021)
13. Hu, H.H., Yokoo, T., Bashir, M.R., Sirlin, C.B., Hernando, D., Malyarenko, D., Chenevert, T.L., Smith, M.A., Serai, S.D., Middleton, M.S., et al.: Linearity and bias of proton density fat fraction as a quantitative imaging biomarker: a multicenter, multiplatform, multivendor phantom study. Radiology **298**(3), 640 (2021)
14. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. Nature methods **18**(2), 203–211 (2021)
15. Joskowicz, L., Cohen, D., Caplan, N., Sosna, J.: Inter-observer variability of manual contour delineation of structures in ct. European radiology **29**(3), 1391–1399 (2019)
16. Kerfoot, E., Clough, J., Oksuz, I., Lee, J., King, A.P., Schnabel, J.A.: Left-ventricle quantification using residual u-net. In: International Workshop on Statistical Atlases and Computational Models of the Heart. pp. 371–380. Springer (2018)
17. Kway, Y.M., Thirumurugan, K., Tint, M.T., Michael, N., Shek, L.P.C., Yap, F.K.P., Tan, K.H., Godfrey, K.M., Chong, Y.S., Fortier, M.V., et al.: Automated segmentation of visceral, deep subcutaneous, and superficial subcutaneous adipose tissue volumes in mri of neonates and young children. Radiology: Artificial Intelligence **3**(5) (2021)
18. Larciprete, G., Valensise, H., Di Pierro, G., Vasapollo, B., Casalino, B., Arduini, D., Jarvis, S., Cirese, E.: Intrauterine growth restriction and fetal body composition. Ultrasound Obstetr Gynecol **26**(3), 258–262 (2005)
19. Lee, W., Balasubramaniam, M., Deter, R.L., Yeo, L., Hassan, S.S., Gotsch, F., Kusanovic, J.P., Gonçalves, L.F., Romero, R.: New fetal weight estimation models using fractional limb volume. Ultrasound Obstetr Gynecol **34**(5), 556–565 (2009)

20. Lee, W., Deter, R.L., McNie, B., Gonçalves, L.F., Espinoza, J., Chaiworapongsa, T., Balasubramaniam, M., Romero, R.: The fetal arm: individualized growth assessment in normal pregnancies. J Ultrasound Med **24**(6), 817–828 (2005)
21. Lin, D., Wang, Z., Li, H., Zhang, H., Deng, L., Ren, H., Sun, S., Zheng, F., Zhou, J., Wang, M.: Automated measurement of pancreatic fat deposition on dixon mri using nnu-net. J Magn Reson Imaging (2022)
22. Mack, L.M., Kim, S.Y., Lee, S., Sangi-Haghpeykar, H., Lee, W.: A novel semiautomated fractional limb volume tool for rapid and reproducible fetal soft tissue assessment. J Ultrasound Med **35**(7), 1573–1578 (2016)
23. Meshaka, R., Gaunt, T., Shelmerdine, S.C.: Artificial intelligence applied to fetal mri: A scoping review of current research. Br. J. Radiol. **95**, 20211205 (2022)
24. Roelants, J., Vermeulen, M., Koning, I., Groenenberg, I., Willemsen, S., Hokken-Koelega, A., Joosten, K., Reiss, I., Steegers-Theunissen, R.: Foetal fractional thigh volume: an early 3d ultrasound marker of neonatal adiposity. Pediatric obesity **12**, 65–71 (2017)
25. Nobile de Santis, M., Taricco, E., Radaelli, T., Spada, E., Rigano, S., Ferrazzi, E., Milani, S., Cetin, I.: Growth of fetal lean mass and fetal fat mass in gestational diabetes. Ultrasound Obstetr Gynecol **36**(3), 328–337 (2010)
26. Shamshad, F., Khan, S., Zamir, S.W., Khan, M.H., Hayat, M., Khan, F.S., Fu, H.: Transformers in medical imaging: A survey. arXiv preprint arXiv:2201.09873 (2022)
27. Shaw, M., Lutz, T., Gordon, A.: Does low body fat percentage in neonates greater than the 5th percentile birthweight increase the risk of hypoglycaemia and neonatal morbidity? Journal of Paediatrics and Child Health **55**(12), 1424–1428 (2019)
28. Tang, Y., Yang, D., Li, W., Roth, H.R., Landman, B., Xu, D., Nath, V., Hatamizadeh, A.: Self-supervised pre-training of swin transformers for 3d medical image analysis. In: Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. pp. 20730–20740 (2022)
29. Torrents-Barrena, J., Piella, G., Masoller, N., Gratacós, E., Eixarch, E., Ceresa, M., Ballester, M.Á.G.: Segmentation and classification in mri and us fetal imaging: recent trends and future prospects. Medical Image Analysis **51**, 61–88 (2019)
30. Yushkevich, P.A., Gao, Y., Gerig, G.: Itk-snap: an interactive tool for semiautomatic segmentation of multi-modality biomedical images. In: Int Conf of IEEE Engineering in Medicine and Biology Society (2016)